



FortiOS™ Handbook - Firewall

VERSION 5.2.0

FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

FORTIGATE COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING SERVICES

<http://www.fortinet.com/training>

FORTIGUARD CENTER

<http://www.fortiguard.com>

FORTICAST

<http://forticast.fortinet.com>

END USER LICENSE AGREEMENT

<http://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdocs@fortinet.com



Friday, May 5, 2017

FortiOS™ Handbook - Firewall

01-520-1248222-20140707

TABLE OF CONTENTS

Change log	8
Introduction	9
How this guide is organized	9
FortiGate Firewall Components	10
How does a FortiGate protect your network?	11
Firewall concepts	13
What is a firewall?	13
Network layer or packet filter firewalls	13
Application layer firewalls	14
Proxy servers	14
Security profiles	15
IPv6	16
The origins of NAT	17
Static NAT	18
Dynamic NAT	18
Benefits of NAT	19
NAT in transparent mode	21
Central NAT table	21
IPv6 in FortiOS	22
Dual stack routing configuration	23
IPv6 tunneling	23
Tunneling IPv6 through IPsec VPN	24
NAT	24
NAT64 and NAT46	24
NAT66	25
How FortiOS differentiates sessions when NATing	26
How packets are handled by FortiOS	32
FortiGate modes	34
NAT/Route mode	34
Transparent mode	34
Quality of Service	34
Traffic policing	35
Traffic shaping	35
Queuing	35

Interfaces and zones.....	35
Firewall objects.....	37
Addresses.....	37
Interfaces.....	38
IPv4 addresses.....	38
Subnet addresses.....	39
IP range addresses.....	40
FQDN addresses.....	41
Geography based addresses.....	43
URL pattern addresses.....	45
Address groups.....	46
UUID support.....	46
Virtual IPs.....	48
Dynamic VIP according to DNS translation.....	49
Creating a virtual IP.....	49
Virtual IP groups.....	52
Creating a virtual IP group.....	52
IP pools.....	53
Creating a dynamic IP pool.....	55
Fixed port.....	55
Match-VIP.....	55
Services and TCP ports.....	56
Categories.....	56
Protocol types.....	57
Service groups.....	75
Example scenario: Supporting audio/visual conferencing.....	76
VIP.....	76
Creating an address for the subnet.....	77
Configuring the services.....	77
Creating the service group.....	80
Creating the IPS security profile.....	80
Policies.....	81
Firewall schedules.....	84
Schedule expiration.....	84
Firewall-session-dirty setting.....	85
Schedule groups.....	86
Schedule expiration.....	86
Security policies.....	87
Firewall policies.....	87
Firewall policy parameters.....	88
What is not expressly allowed is denied.....	90

Policy order.....	90
Policy identification.....	93
UUID support.....	93
Security profiles.....	94
AntiVirus.....	94
Web filtering.....	95
Application control.....	95
Intrusion protection (IPS).....	95
Email filtering.....	95
Data Leak Prevention (DLP).....	96
VoIP.....	96
ICAP.....	96
EndPoint control.....	96
Proxy option components.....	96
The use of different proxy profiles and profile options.....	97
SSL/SSH inspection.....	99
Inspection exemption.....	100
Creating or editing an SSL/SSH inspection profile.....	101
Viewing firewall policies.....	103
How “Any” policy can remove the section view.....	104
Security policy configuration extensions.....	104
Identity-based policies.....	105
Identity-based policy positioning.....	105
VPN policies.....	107
IPsec policies.....	107
Interface policies.....	107
DoS protection.....	108
One-arm IDS.....	112
IPv6 IPS.....	112
Traffic destined to the FortiGate unit.....	112
Dropped, flooded, broadcast, multicast and L2 packets.....	112
GUI and CLI.....	113
Local-In policies.....	113
Security policy 0.....	114
Deny policies.....	114
Accept policies.....	115
Fixed port.....	115
Endpoint security.....	115
Traffic logging.....	116
Quality of service.....	117
Queuing.....	117
Policy Monitor.....	118

Upper pane.....	118
Lower pane.....	118
Network defense.....	120
Monitoring.....	120
Blocking external probes.....	120
Address sweeps.....	120
Port scans.....	121
Probes using IP traffic options.....	121
Evasion techniques.....	122
Defending against DoS attacks.....	125
The “three-way handshake”.....	125
SYN flood.....	125
SYN spoofing.....	126
DDoS SYN flood.....	127
Configuring the SYN threshold to prevent SYN floods.....	127
SYN proxy.....	127
Other flood types.....	128
DoS policies.....	128
GUI and CLI - what you may not know.....	129
Mouse tricks.....	130
Changing the default column setting on the policy page.....	131
Example.....	132
Naming rules and restrictions.....	133
Character restrictions.....	134
Numeric values.....	135
Numeric values.....	136
Selecting options from a list.....	137
Enabling or disabling options.....	138
To enable or disable optionally displayed features.....	140
Building firewall objects and policies.....	141
Building firewall objects and policies.....	142
Example: IPv4 firewall addresses.....	142
Scenario: Mail server.....	142
Scenario: First floor network.....	143
Scenario: Marketing department.....	144
Verification.....	144
Example: IPv6 firewall addresses.....	145
Scenario: Mail server.....	145
Scenario: First floor network.....	145
Verification.....	146
Example: FQDN address.....	146
Verification.....	147

Changing the TTL of a FQDN address	147
Example: Geography-based address	147
Example: IPv4 address group	148
Example: IPv6 address group	149
Example: Multicast address	150
Example: Service category	151
Example: TCP/UDP/SCTP service	152
Example: ICMP service	154
Example: ICMPv6 service	155
Example: Service group	156
Example: Virtual IP address	157
VIP group	158
Example: IP pool	159
Example: Central NAT table	160
Example: Firewall schedule - recurring	161
Example: Firewall schedule - one-time	162
Example: Schedule group	163
Example: Proxy option	164
Oversized files	166
Example: DoS policy	167
Multicast forwarding	171
Sparse mode	171
Dense mode	173
Multicast IP addresses	173
PIM Support	174
Multicast forwarding and FortiGate units	175
Multicast forwarding and RIPv2	175
Configuring FortiGate multicast forwarding	176
Adding multicast security policies	177
Enabling multicast forwarding	177
Multicast routing examples	180
Example FortiGate PIM-SM configuration using a static RP	180
FortiGate PIM-SM debugging examples	189
Example multicast destination NAT (DNAT) configuration	193
Example PIM configuration that uses BSR to find the RP	196

Change log

Date	Change Description
2017-05-3	Additional content added to DoS Protection
2016-01-29	How FortiOS differentiates sessions when NATing
2015-09-10	added content on log-invalid-packet setting
2014-06-12	Added content regarding VIP and polices in policy order exception, Deny policies and VIP objects.
2014-12-30	Resolved bug 0257589 (wild card values in FQDN address objects)
2014-12-22	Resolved bug 0256490 (srcaddr and dstaddr objects in multicast policies)
2014-12-19	Cross reference links removed
2014-07-07	Initial release.

Introduction

Welcome and thank you for selecting Fortinet products for your network protection. This document is intended to provide the concepts and techniques that will be needed to configure the FortiGate firewall on your FortiGate unit.

Before you start administrating your FortiGate device, certain assumptions have been made in the writing of this manual:

- You have administrative access to the Web-based GUI or to the Command Line Interface.
- The FortiGate unit is integrated into your network.
- The operation mode (NAT or Transparent) has been configured.
- Network interfaces have been configured.
- DNS settings have been configured.
- The system time settings have been configured.
- Firmware is up to date.
- FortiGuard Service licenses are current and the device is able to connect to the FortiGuard servers.
- If you are using FortiCloud, it is properly configured.

How this guide is organized

"Firewall concepts" explains the ideas behind the components, techniques and processes that are involved in setting up and running a firewall in general, and the FortiGate firewall in particular. Regardless of how experienced someone is with firewalls, as they go through the process of configuring a firewall new to them they are likely to come across a term or setting that they may not be familiar with. FortiGate firewalls are quite comprehensive and can be very granular in the functions they perform, so it makes sense to have a consistent frame of reference for the ideas we will be working with.

Some examples of the concepts that will be addressed here are:

- "What is a Firewall?"
- "NAT"
- "IPv6"

"Firewall objects" describes the following:

- Addressing
- Services
- Firewall Policies

"Network defense" describes various methods of defending your network using the abilities of the FortiGate firewall.

"GUI and CLI - what you may not know" helps you navigate and find the components in the Web-based Manager that you will need to build the functions. This section does not include any in-depth explanations of what each object does as that is covered in the concepts section. This section is for showing you where you need to input your information and letting you know what format the interface expects to get that information

"Building firewall objects and policies" is similar to a cookbook in that it will refer to a number of common tasks that you will likely perform to get the full functionality out of your FortiGate firewall. Because of the way that firewalls are designed, performing many of the tasks requires that firewall components be set up in a number of different sections of the interface and be configured to work together to achieve the desired result. This section will bring those components together as a straight forward series of instructions.

"Multicast forwarding" is a reference guide including the concepts and examples that are involved in the use of multicast addressing and policy forwarding as it is used in the FortiGate firewall.

FortiGate Firewall Components

The FortiGate firewall is made up of a number of different components that are used to build an impressive list of features. These features have flexibility of scope and granularity of control that provide protection far beyond that provided by the basic firewalls of the past.

Some of the components that FortiOS uses to build features are:

- Interfaces
- VLANs
- Soft switches
- Zones
- Predefined addresses
 - IP address-based
 - FQDN-based
 - Geography-based
- Access schedules
- Authentication
 - Local user-based
 - Authentication server-based (Active Directory, RADIUS, LDAP)
 - Device-based
- Configurable services
- IPv4 and IPv6 protocol support

The features of FortiOS include but are not limited to:

- Security profiles, sometimes referred to as Unified Threat Management (UTM) or Next Generation Firewall (NGFW)
- Predefined firewall addresses (this includes IPv4 and IPv6, IP pools, wildcard addresses and netmasks, and geography-based addresses)
- Monitoring traffic
- Traffic shaping and per-IP traffic shaping (advanced)
- Firewall schedules
- Services (such as AOL, DHCP and FTP)
- Logging traffic
- Quality of Service (QoS)
- Identity-based policies
- Endpoint security

The "Firewall concepts" expand on what each of the features does and how they relate to the administration of the FortiGate firewall. The section will also try to explain some of the common firewall concepts that will be touched on in the implementing of these features.

"Building firewall objects and policies" shows how to perform specific tasks with the FortiGate firewall.

How does a FortiGate protect your network?

The FortiGate firewall protects your network by taking the various components and using them together to build a kind of wall or access control point so anyone that is not supposed to be on your network is prevented from accessing your network in any way other than those approved by you. It also protects your network from itself by keeping things that shouldn't happen from happening and optimizing the flow of traffic so the network is protected from traffic congestion that would otherwise impede traffic flow.

Most people have at one time or another played with a child's toy system made up of interlocking blocks. The blocks come in different shapes and sizes so you can build structures to suit your needs. The components of the FortiGate firewall are similar. You are not forced to use all of the blocks all of the time. You mix and match them to get the results that you are looking for. You can build a very basic structure, where its only function is to direct traffic in and out to the correct subnets. You can build a fortress that only allows specific traffic to or from specific hosts at specific times of day and only when credentials that have been pre-approved have been provided. You can also add in that all of the traffic is encrypted, so that even when the traffic is out on the Internet it is private from the world. Just like the interlocking blocks, what you build is up to you, but chances are if you put them together the right way there isn't much that can't be built.

Here is one example of how the components could be put together to support the requirements of a network infrastructure design.

- Off the internal interface you could have separate VLANs. One for each department of Sales, Marketing and Engineering so the traffic from the users on one VLAN does not intrude upon the hosts of the other VLANs and each department is isolated from the others for security reasons.
- To ease in the administration, each of the VLAN sub-interfaces is made a member of a zone so security policies that apply to all hosts on all the VLANs can be applied to all of them at once.
- Using the addresses component, each of the IP address ranges could be assigned a user-friendly name so they could be referred to individually. For policies that would refer to them all as a whole, the individual ranges can be made members of an address group.
- Firewall schedules could be created to address the differing needs of each of the groups so that Sales and Marketing could be allowed access to the Internet during regular business hours and the Engineering department could be allowed access during the lunch break.
- By setting up the outgoing policies to use FortiGuard Web-filtering, the employees could be prevented from visiting inappropriate sites and thus enforcing the policies of the HR department.
- A couple of virtual IP addresses with port forwarding could be configured to allow users on the Internet access to a web server on the DMZ subnet using the company's only Public IP address. This would not affect the traffic that goes to the company's mail server hosted on a completely different computer.
- Even though the Web server on the same DMZ has an FTP service to allow for the uploading of web pages to the web server from the Marketing and Engineer teams, by placing a DENY policy on any FTP traffic from the Internet, malicious users are prevented from abusing the FTP service.
- By monitoring the traffic as it goes through the policies, you can verify that the policies are in working order.
- By using a combination of ALLOW and DENY policies and placing them in the correct order, you could arrange for an outside contractor to be allowed to update the web site as well.

This list of possible configurations is not extensive, but it does give an idea of how different components can be mixed and matched to build a configuration that meets an organization's needs and at the same time protect it from security risks.

Firewall concepts

There are a number of concepts that are consistent throughout the firewall industry. Having a solid grasp of these ideas and terms can give you a better idea of what your FortiGate firewall is capable of and how it will fit within your network's architecture.

This chapter describes the following firewall concepts:

- What is a Firewall?
- IPv6
- NAT
- How packets are handled by FortiOS
- FortiGate modes
- Quality of Service
- Interfaces and zones

What is a firewall?

The term firewall originally referred to a wall intended to confine a fire or potential fire within a building. Later uses refer to similar structures, such as the metal sheet separating the engine compartment of a vehicle or aircraft from the passenger compartment.

A firewall can either be software-based or hardware-based and is used to help keep a network secure. Its primary objective is to control the incoming and outgoing network traffic by analyzing the data packets and determining whether it should be allowed through or not, based on a predetermined rule set. A network's firewall builds a bridge between an internal network that is assumed to be secure and trusted, and another network, usually an external (inter)network such as the Internet, that is not assumed to be secure and trusted.

Network layer or packet filter firewalls

Stateless firewalls

Stateless firewalls are the oldest form of these firewalls. They are faster and simple in design requiring less memory because they process each packet individually and don't require the resources necessary to hold onto packets like stateful firewalls. Stateful firewalls inspect each packet individually and check to see if it matches a predetermined set of rules. According to the matching rule, the packet is either allowed, dropped or rejected. In the case of a rejection, an error message is sent to the source of the traffic. Each packet is inspected in isolation and information is only gathered from the packet itself. Simply put, if the packets were not specifically allowed according to the list of rules held by the firewall, they were not getting through.

Stateful firewalls

Stateful firewalls retain packets in memory so that they can maintain context about active sessions and make judgments about the state of an incoming packet's connection. This enables Stateful firewalls to determine if a packet is the start of a new connection, a part of an existing connection, or not part of any connection. If a packet

is part of an existing connection based on comparison with the firewall's state table, it will be allowed to pass without further processing. If a packet does not match an existing connection, it will be evaluated according to the rules set for new connections. Predetermined rules are used in the same way as a stateless firewall but they can now work with the additional criteria of the state of the connection to the firewall.

Best Practices Tip for improving performance:



Blocking the packets in a denied session can take more cpu processing resources than passing the traffic through. By putting denied sessions in the session table, they can be kept track of in the same way that allowed sessions are so that the FortiGate unit does not have to redetermine whether or not to deny all of the packets of a session individually. If the session is denied all packets of that session are also denied.

In order to configure this you will need to use 2 CLI commands

```
config system setting
    set ses-denied-traffic enable
    set block-session-timer <integer 1 - 300> (this determines in
seconds, how long, the session is kept in the table)
end
```

Application layer firewalls

Application layer filtering is yet another approach and as the name implies it works primarily on the Application Layer of the OSI Model.

Application Layer Firewalls actually, for lack of a better term, understand certain applications and protocols. Examples would be FTP, DNS and HTTP. This form of filtration is able to check to see if the packets are actually behaving incorrectly or if the packets have been incorrectly formatted for the protocol that is indicated. This process also allows for the use of deep packet inspection and the sharing of functionality with Intrusion Prevention Systems (IPS).

Application-layer firewalls work on the application level of the TCP/IP stack (i.e., all browser traffic, or all telnet or ftp traffic), and may intercept all packets traveling to or from an application. They block other packets (usually dropping them without acknowledgment to the sender). Application firewalls work much like a packet filter but application filters apply filtering rules (allow/block) on a per process basis instead of filtering connections on a per port basis.

On inspecting all packets for improper content, firewalls can restrict or prevent outright the spread of networked computer worms and trojans. The additional inspection criteria can add extra latency to the forwarding of packets to their destination.

Proxy servers

A proxy server is an appliance or application that acts as an intermediary for communicating between computers. A computer has a request for information. The packets are sent to the designated resource but before they can get there they are blocked by the proxy server saying that it will take the request and pass it on. The Proxy Server processes the request and if it is valid it passes onto the designated computer. The designated computer gets the packet and processes the request, sending the answer back to the proxy server. The proxy server sends the information back to the originating computer. It's all a little like a situation with two people who refuse to talk directly with each other using someone else to take messages back and forth.

From a security stand point a Proxy Server can serve a few purposes:

- Protects the anonymity of the originating computer
- The two computers never deal directly with each other
- Packets that are not configured to be forwarded are dropped before reaching the destination computer.
- If malicious code is sent it will affect the Proxy server with out affecting the originating or sending computer.

Proxies can perform a number of roles including:

- Content Filtering
- Caching
- DNS proxy
- Bypassing Filters and Censorship
- Logging and eavesdropping
- Gateways to private networks
- Accessing service anonymously

Security profiles

Unified Threat Management and Next Generation Firewall are terms originally coined by market research firms and refer to the concept of a comprehensive security solution provided in a single package. It is basically combining of what used to be accomplished by a number of different security technologies all under a single umbrella or in this case, a single device. On the FortiGate firewall this is achieved by the use of Security Profiles and optimized hardware.

In effect it is going from a previous style of firewall that included among its features:

- Gateway Network Firewall
- Routing
- VPN

To a more complete system that includes:

- Gateway Network Firewall
- Routing
- VPN
- Traffic Optimization
- Proxy Services
- Content Filtering
- Application Control
- Intrusion Protection
- Denial of Service Attack Protection
- Anti-virus
- Anti-spam
- Data Leak Prevention
- Endpoint Control of Security Applications
- Load Balancing
- WiFi Access Management
- Authentication Integration into Gateway Security

- Logging
- Reporting

Advantages of using security profiles

- Avoidance of multiple installations.
- Hardware requirements are fewer.
- Fewer hardware maintenance requirements.
- Less space required.
- Compatibility - multiple installations of products increase the probability of incompatibility between systems.
- Easier support and management.
- There is only one product to learn therefore a reduced requirement of technical knowledge.
- Only a single vendor so there are fewer support contracts and Service Level Agreements.
- Easier to incorporate into existing security architecture.
- Plug and play architecture.
- Web based GUI for administration.

IPv6

Internet Protocol version 6 (IPv6) will succeed IPv4 as the standard networking protocol of the Internet. IPv6 provides a number of advances over IPv4 but the primary reason for its replacing IPv4 is its limitation in addresses. IPv4 uses 32 bit addresses which means there is a theoretical limit of 2 to the power of 32. The IPv6 address scheme is based on a 128 bit address or a theoretical limit of 2 to the power of 128.

Possible Addresses:

- IPv4 = 4,294,967,296 (over 4 billion)
- IPv6 = 340,282,366,920,938,463,463,374,607,431,768,211,456 (over 340 undecillion - We had to look that term up. We didn't know what a number followed by 36 digits was either)

Assuming a world population of approximately 8 billion people, IPv6 would allow for each individual to have approximately 42,535,295,865,117,200,000,000,000,000 devices with an IP address. That's 42 quintillion devices.

There is little likelihood that you will ever need to worry about these numbers as any kind of serious limitation in addressing but they do give an idea of the scope of the difference in the available addressing.

Aside from the difference of possible addresses there is also the different formatting of the addresses that will need to be addressed.

A computer would view an IPv4 address as a 32 bit string of binary digits made up of 1s and 0s, broken up into 4 octets of 8 digits separated by a period "."

Example:

10101100.00010000.11111110.00000001

To make number more user friendly for humans we translate this into decimal, again 4 octets separated by a period "." which works out to:

172.16.254.1

A computer would view an IPv6 address as a 128 bit string of binary digits made up of 1s and 0s, broken up into 8 octets of 16 digits separated by a colon “:”

```
1000000000000001:0000110110111000:101011000001000:1111111000000001:0000000000000000  
0:0000000000000000:0000000000000000:0000000000000000
```

To make number a little more user friendly for humans we translate this into hexadecimal, again 8 octets separated by a colon “:” which works out to:

```
8001:0DB8:AC10:FE01:0000:0000:0000:0000:
```

Because any four-digit group of zeros within an IPv6 address may be reduced to a single zero or altogether omitted, this address can be shortened further to:

```
8001:0DB8:AC10:FE01:0:0:0:0
```

or

```
8001:0DB8:AC10:FE01::
```

Some of the other benefits of IPv6 include:

- More efficient routing
- Reduced management requirement
- Stateless auto-reconfiguration of hosts
- Improved methods to change Internet Service Providers
- Better mobility support
- Multi-homing
- Security
- Scoped address: link-local, site-local and global address space

The origins of NAT

In order to understand NAT it helps to know why it was created. At one time, every computer that was part of a network had to have it's own addresses so that the other computers could talk to it. There were a few protocols in use at the time, some of which were only for use on a single network, but of those that were routable, the one that had become the standard for the Internet was IP (Internet Protocol) version 4.

When IP version 4 addressing was created nobody had any idea how many addresses would be needed. The total address range was based on the concept of 2 to the 32nd power, which works out to be 4 294 967 296 potential addresses. Once you eliminate some of those for reserved addresses, broadcast addresses, network addresses, multicasting, etc., you end up with a workable scope of about 3.2 million addressees. This was thought to be more than enough at the time. The designers were not expecting the explosion of personal computing, the World Wide Web or smart phones. As of the beginning of 2012, some estimate the number of computers in the world in the neighborhood of 1 billion, and most of those computer users are going to want to be on the Internet or Search the World Wide Web. In short, we ran out of addresses.

This problem of an address shortage was realized before we actually ran out, and in the mid 1990s 2 technical papers called RFCs numbered 1631 (<http://www.ietf.org/rfc/rfc1631.txt>) and 1918 (<http://tools.ietf.org/html/rfc1918>), proposed components of a method that would be used as a solution until a new addressing methodology could be implemented across the Internet infrastructure. For more information on this you can look up IP version 6.

RFC 1631 described a process that would allow networking devices to translate a single public address to multiple private IP addresses and RFC 1918 laid out the use of the private addresses. The addresses that were on the Internet (Public IP addresses) could not be duplicated for them to work as unique addresses, but behind a firewall, which most large institutions had, they could use their own Private IP addresses for internal use and the internal computers could share the external or Public IP address.

To give an idea on a small scale how this works, imagine that a company has a need for 200 computer addresses. Before Private IP addresses and NAT the company would have purchased a full Class C address range which would have been 254 usable IP addresses; wasting about 50 addresses. Now with NAT, that company only needs 1 IP address for its 200 computers and this leaves the rest of the IP addresses in that range available for other companies to do the same thing.

NAT gives better value than it would first appear because it is not 253 companies that can use 254 addresses but each of those 254 companies could set up their networking infrastructures to use up to thousands of Private IP addresses, more if they don't all have to talk to the Internet at the same time. This process enabled the Internet to keep growing even though we technically have many more computers networked than we have addresses.

Static NAT

In Static NAT one internal IP address is always mapped to the same public IP address.

In FortiGate firewall configurations this is most commonly done with the use of Virtual IP addressing.

An example would be if you had a small range of IP addresses assigned to you by your ISP and you wished to use one of those IP address exclusively for a particular server such as an email server.

Say the internal address of the Email server was 192.168.12.25 and the Public IP address from your assigned addresses range from 256.16.32.65 to 256.16.32.127. Many readers will notice that because one of the numbers is above 255 that this is not a real Public IP address. The Address that you have assigned to the interface connected to your ISP is 256.16.32.66, with 256.16.32.65 being the remote gateway. You wish to use the address of 256.16.32.70 exclusively for your email server.

When using a Virtual IP address you set the external IP address of 256.16.32.70 to map to 192.168.12.25. This means that any traffic being sent to the public address of 256.16.32.70 will be directed to the internal computer at the address of 192.168.12.25

When using a Virtual IP address, this will have the added function that when ever traffic goes from 192.168.12.25 to the Internet it will appear to the recipient of that traffic at the other end as coming from 256.16.32.70.

You should note that if you use Virtual IP addressing with the Port Forwarding enabled you do not get this reciprocal effect and must use IP pools to make sure that the outbound traffic uses the specified IP address.

Dynamic NAT

Dynamic NAT maps the private IP addresses to the first available Public Address from a pool of possible Addresses. In the FortiGate firewall this can be done by using IP Pools.

Overloading

This is a form of Dynamic NAT that maps multiple private IP address to a single Public IP address but differentiates them by using a different port assignment. This is probably the most widely used version of NAT. This is also referred to as PAT (Port Address Translation) or Masquerading.

An example would be if you had a single IP address assigned to you by your ISP but had 50 or 60 computers on your local network.

Say the internal address of the interface connected to the ISP was 256.16.32.65 (again an impossible address) with 256.16.32.64 being the remote gateway. If you are using this form of NAT any time one of your computers accesses the Internet it will be seen from the Internet as 256.16.32.65. If you wish to test this go to 2 different computers and verify that they each have a different private IP address then go to a site that tells you your IP address such as www.ipchicken.com. You will see that the site gives the same result of 256.16.32.65, if it existed, as the public address for both computers.

As mentioned before this is sometimes called Port Address Translation because network device uses TCP ports to determine which internal IP address is associated with each session through the network device. For example, if you have a network with internal addresses ranging from 192.168.1.1 to 192.168.1.255 and you have 5 computers all trying to connect to a web site which is normally listening on port 80 all of them will appear to the remote web site to have the IP address of 256.16.32.65 but they will each have a different sending TCP port, with the port numbers being somewhere between 1 and 65 535, although the port numbers between 1 to 1024 are usually reserved or already in use. So it could be something like the following:

192.168.1.10	256.16.32.65:	port 486
192.168.1.23	256.16.32.65:	port 2409
192.168.1.56	256.16.32.65:	port 53763
192.168.1.109	256.16.32.65:	port 5548
192.168.1.201	256.16.32.65:	port 4396

And the remote web server would send the responding traffic back based on those port numbers so the network device would be able to sort through the incoming traffic and pass it on to the correct computer.

Overlapping

Because everybody is using the relative same small selection of Private IP addresses it is inevitable that there will be two networks that share the same network range that will need to talk with each other. This happens most often over Virtual Private Networks or when one organization ends up merging with another. This is a case where a private IP address may be translated into a different private IP address so there are no issues with conflict of addresses or confusion in terms of routing.

An example of this would be when you have a Main office that is using an IP range of 172.16.0.1 to 172.20.255.255 connecting through a VPN to a recently acquired branch office that is already running with an IP range of 172.17.1.1 to 172.17.255.255. Both of these ranges are perfectly valid but because the Branch office range is included in the Main Office range any time the system from the Main office try to connect to an address in the Branch Office the routing the system will not send the packet to the default gateway because according to the routing table the address is in its own subnet.

The plan here would be to NAT in both directions so that traffic from neither side of the firewall would be in conflict and they would be able to route the traffic. Everything coming from the Branch Office could be assigned an address in the 192.168.1.1 to 192.168.1.255 range and everything from the Main office going to the Branch Office could be assigned to an address in the 192.168.10.1 to 192.168.10.255 range.

Benefits of NAT

More IP addresses available while conserving public IP addresses

As explained earlier, this was the original intent of the technology and does not need to be gone into further.

Financial savings

Because an organization does not have to purchase IP addresses for every computer in use there is a significant cost savings due to using the process of Network Address Translation.

Security enhancements

One of the side benefits of the process of NAT is an improvement in security. Individual computers are harder to target from the outside and if port forwarding is being used computers on the inside of a firewall are less likely to have unmonitored open ports accessible from the Internet.

Ease of compartmentalization of your network

With a large available pool of IP addresses to use internally a network administrator can arrange things to be compartmentalized in a rational and easily remembered fashion and networks can be broken apart easily to isolate for reasons of network performance and security.

Example

You have a large organization that for security reasons has certain departments that do not share network resources.

You can have the main section of the organization set up as follows;

Network Devices	192.168.1.1 to 192.168.1.25
Internal Servers	192.168.1.26 to 192.168.1.50
Printers	192.168.1.51 to 192.168.1.75
Administration Personnel	192.168.1.76 to 192.168.1.100
Sales People	192.168.1.101 to 192.168.1.200
Marketing	192.168.1.201 to 192.168.1.250

You could then have the following groups broken off into separate subnets:

Accounting	192.168.100.1 to 192.168.100.255
Research and Development	172.16.1.1 to 172.16.255.255
Executive Management	192.168.50.1 to 192.168.50.255
Web sites and Email Servers	10.0.50.1 to 10.0.50.255

These addresses do not have to be assigned right away but can be used as planned ranges.

NAT in transparent mode

Similar to operating in NAT mode, when operating a FortiGate unit in Transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.
- Add IP pools as required for source address translation

A FortiGate unit operating in Transparent mode normally has only one IP address - the management IP. To support NAT in Transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

Use the following steps to configure NAT in Transparent mode:

1. Add two management IPs
2. Add an IP pool to the WAN1 interface
3. Add an Internal to WAN1 security policy

You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

The usual practice of NATing in transparent mode makes use of two management IP addresses that are on different subnets, but this is not an essential requirement in every case.

If there is a router between the client systems and the FortiGate unit you can use the router's capabilities of tracking sessions to assign NATed addresses from an IP pool to the clients even if the assigned address don't belong to a subnet on your network.

Example

Client computer has an IP address of 1.1.1.33 on the subnet 1.1.1.0/24.

Router "A" sits between the client computer and the FortiGate (in Transparent mode) with the IP address of 1.1.1.1 on the client's side of the router and the IP address of 192.168.1.211 on the FortiGate's side of the router.

Use NAT to assign addresses from an address pool of 9.9.9.1 to 9.9.9.99 to traffic coming from gateway of 192.168.1.211.

To enable the return traffic to get to the original computer, set up a static route that assigns any traffic with a destination of 9.9.9.0/24 to go through the 192.168.1.211 gateway. As long as the session for the outgoing traffic has been maintained, communication between the client computer and the external system on the other side of the FortiGate will work.

Central NAT table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. The NAT policies can be rearranged within the policy list as well. NAT policies are applied to network traffic after a security policy.

IPv6 in FortiOS

From an administrative point of view IPv6 works almost the same as IPv4 in FortiOS. The primary difference is the use IPv6 format for addresses. There is also no need for NAT if the FortiGate firewall is the interface between IPv6 networks. If the subnets attached to the FortiGate firewall are IPv6 and IPv4 NAT can be configured between the 2 different formats. This will involve either configuring a dual stack routing or IPv4 tunneling configuration. The reason for this is simple. NAT was developed primarily for the purpose of extending the number of usable IPv4 addresses. IPv6's addressing allows for enough available addresses so the NAT is no longer necessary.

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network. Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

By default the IPv6 settings are not displayed in the Web-based Manager. It is just a matter of enabling the display of these feature to use them through the web interface. To enable them just go to **System > Admin > Settings** and select **IPv6 Support on GUI**. Once enabled, you will be able to use IPv6 addresses as well as the IPv4 addressing for the following FortiGate firewall features:

- Static routing
- Policy Routing
- Packet and network sniffing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- IPsec VPN
- DNS
- DHCP
- SSL VPN
- Network interface addressing
- Security Profiles protection
- Routing access lists and prefix lists
- NAT/Route and Transparent mode
- NAT 64 and NAT 66
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Logging and reporting
- Security policies

- SNMP
- Authentication
- Virtual IPs and groups
- IPv6 over SCTP
- IPv6-specific troubleshooting, such as ping6

Dual stack routing configuration

Dual stack routing implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. A dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary. Devices that are on this type of network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses. In the FortiOS dual stack architecture it is not just the basic addressing functions that operate in both versions of IP. The other features of the appliance such as Security Profiles and routing can also use both IP stacks.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunneling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

IPv6 tunneling

IPv6 Tunneling is the act of tunneling IPv6 packets from an IPv6 network through an IPv4 network to another IPv6 network. This is different than Network Address Translation (NAT) because once the packet reaches its final destination the true originating address of the sender will still be readable. The IPv6 packets are encapsulated within packets with IPv4 headers, which carry their IPv6 payload through the IPv4 network. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

The key to IPv6 tunneling is the ability of the 2 devices, whether they are a host or a network device, to be dual stack compatible. They have to be able to work with both IPv4 and IPv6 at the same time. In the process the entry node of the tunnel portion of the path will create an encapsulating IPv4 header and transmit the encapsulated packet. The exit node at the end of the tunnel receives the encapsulated packet. The IPv4 header is removed. The IPv6 header is updated and the IPv6 packet is processed.

There are two types of tunnels in IPv6:

Automatic tunnels	Automatic tunnels are configured by using IPv4 address information embedded in an IPv6 address – the IPv6 address of the destination host includes information about which IPv4 address the packet should be tunneled to.
Configured tunnels	Configured tunnels must be configured manually. These tunnels are used when using IPv6 addresses that do not have any embedded IPv4 information. The IPv6 and IPv4 addresses of the endpoints of the tunnel must be specified.

Tunnel configurations

There are a few ways in which the tunneling can be performed depending on which segment of the path between the end points of the session the encapsulation takes place.

Network Device to Network Device	Dual stack capable devices connected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans one segment of the path taken by the IPv6 packets.
Host to Network Device	Dual stack capable hosts can tunnel IPv6 packets to an intermediary IPv6 or IPv4 network device that is reachable through an IPv4 infrastructure. This type of tunnel spans the first segment of the path taken by the IPv6 packets.
Host to Host	Dual stack capable hosts that are interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans the entire path taken by the IPv6 packets.
Network Device to Host	Dual stack capable network devices can tunnel IPv6 packets to their final destination IPv6 or IPv4 host. This tunnel spans only the last segment of the path taken by the IPv6 packets.

Regardless of whether the tunnel starts at a host or a network device, the node that does the encapsulation needs to maintain soft state information, such as the maximum transmission unit (MTU), about each tunnel in order to process the IPv6 packets.

Tunneling IPv6 through IPsec VPN

A variation on the tunneling IPv6 through IPv4 is using an IPsec VPN tunnel between two FortiGate devices. FortiOS supports IPv6 over IPsec. In this sort of scenario, 2 networks using IPv6 behind FortiGate units are separated by the Internet, which uses IPv4. An IPsec VPN tunnel is created between the 2 FortiGate units and a tunnel is created over the IPv4 based Internet but the traffic in the tunnel is IPv6. This has the additional advantage of making the traffic secure as well.

NAT

NAT or Network Address Translation is the process that enables a single device such as a router or firewall to act as an agent between the Internet or Public Network and a local or private network. This “agent”, in real time, translates the source IP address of a device on one network interface, usually the Internal, to a different IP address as it leaves another interface, usually the interface connected to the ISP and the Internet. This enables a single public address to represent a significantly larger number of private addresses.

NAT64 and NAT46

NAT64 and NAT46 are the terms used to refer to the mechanism that allows IPv6 addressed hosts to communicate with IPv4 addressed hosts and vice-versa. Without such a mechanism an IPv6 node on a network such as a corporate LAN would not be able to communicate with a web site that was still in a IPv4 only environment and IPv4 environments would not be able to connect to IPv6 networks.

One of these setups involves having at least 2 interfaces, 1 on an IPv4 network and 1 on an IPv6 network. The NAT64 server synthesizes AAAA records, used by IPv6 from A records used by IPv4. This way client-server and peer to peer communications will be able to work between an IPv6 only client and an IPv4 server without making changes to either of the end nodes in the communication transaction. The IPv6 network attached to the FortiGate unit should be a 32 bit segment, (for instance 64:ff9b::/96, see RFC 6052 and RFC 6146). IPv4 address will be embedded into the communications from the IPv6 client.

Because the IPv6 range of addresses is so much larger than the IPv4 range, a one to one mapping is not feasible. Therefore the NAT64 function is required to maintain any IPv6 to IPv4 mappings that it synthesizes. This can be done either statically by the administrator or automatically by the service as the packets from the IPv6 network go through the device. The first method would be a stateless translation and the second would be a stateful translation. NAT64 is designed for communication initiated from IPv6 hosts to IPv4 addresses. It is address mapping like this that allows the reverse to occur between established connections. The stateless or manual method is an appropriate solution when the NAT64 translation is taking place in front of legacy IPv4 servers to allow those specific servers to be accessed by remote IPv6-only clients. The stateful or automatic solution is best used closer to the client side when you have to allow some specific IPv6 clients to talk to any of the IPv4-only servers on the Internet.

There are currently issues with NAT64 not being able to make everything accessible. Examples would be SIP, Skype, MSN, Goggle talk, and sites with IPv4 literals. IPv4 literals being IPv4 addresses that are imbedded into content rather than a FQDN.

Policies that employ NAT64 or NAT46 can be configured from the web-based manager as long as the feature is enabled using the Features setting found at **System > Config > Features**.

- To create a NAT64 policy go to **Policy > Policy > NAT64 Policy** and select **Create New**.
- To create a NAT46 policy go to **Policy > Policy > NAT46 Policy** and select **Create New**.

The difference between these NAT policies and regular policies is that there is no option to use the security profiles and sensors.

NAT66

NAT 66 is Network Address Translation between 2 IPv6 network. The basic idea behind NAT 66 is no different than the regular NAT between IPv4 networks that we are all used to. The difference are in the mechanics of how it is performed, mainly because of the complexity and size of the addresses that are being dealt with.

In an IPv4 world, the reason for the use of NAT was usually one or a combination of the following 3 reasons:

- Improved security - actual addresses behind NAT are virtually hidden
- Amplification of addresses - hundreds of computers can use as little as a single public IP address
- Internal address stability - there is control of internal addressing. The addresses can stay the same even if Internet Service Providers change.

In these days of security awareness the protective properties of NAT are not something that are not normally depended on by themselves to defend a network and with the vastly enlarged IPv6 address scope there is no longer a need to amplify the available addresses. However, the desire to have internal address control still exists. The most common reason for using NAT66 is likely to be the maintaining of the existing address scheme of the internal network despite changes outside of it. Imagine that you have an internal network of 2000 IP addresses and one day the company changes its ISP and thus the addresses assigned to it. Even if most of the addressing is handled by DHCP, changing the address scheme is going to have an impact on operations.

Addressing stability can be achieved by:

- Keeping the same provider - this would depend on the reason for the change. If the cost of this provider has become too expensive this is unlikely. If the ISP is out of business it becomes impossible.
- Transfer the addresses from the old provider to the new one - There is little motivation for an ISP to do you a favor for not doing business with them.
- Get your own autonomous system number - this can be too expensive for smaller organizations.
- NAT - this is the only one on the list that is in the control of IT.

There are differences between NAT66 and IPv4 NAT. Because there is no shortage of addresses most organizations will be given a /48 network that can be translated into another /48 network. This allows for a one to one translation, no need for port forwarding. This is a good thing because port forwarding is more complicated in IPv6. In fact, NAT66 will actually just be the rewriting of the prefix on the address.

Example

If your current IPv6 address is

```
2001:db8:cafe::/48
```

you could change it to

```
2001:db8:fea7::/48
```

There is an exception to the one to one translation. NAT66 cannot translate internal networks that contain 0xffff in bits 49 through 63 - this is due to the way checksums are calculated in TCP/IP: they use the one's-complement representation of numbers which assigns the value zero to both 0x0000 and 0xffff.

How FortiOS differentiates sessions when NATing

The basics of NAT are fairly simple. Many private addresses get translated into a smaller number of public addresses, often just one. The trick is how the FortiGate keeps track of the return traffic because the web server, or what ever device that was out on the Internet is going to be sending traffic back not to the private address behind the FortiGate but to the IP address of the interface on the public side of the FortiGate.

The way this is done is by making each session unique. Most of the attributes that are available in the network packets cannot be changed without changing where the packet will go but because the source port has to be changed anyway in case two computer on the network used the same source port this is a useful way of making each listing of network attributes a unique combination. As a packet goes through the NAT process FortiOS assigns different source ports for each of the internally initiated sessions and keeping track of which port was used for each device in a database until the session has ended. It then becomes a matter of how the port number is selected.

In a very simple example of an environment using NAT, we will use a fictitious university with a rather large student population. So large in fact that they use a subnet of 10.0.0.0/8 as their subnet for workstation IP addresses. All of these private IP addresses are NATed out a single IP address. To keep the number of numeric values in this example from getting to a confusing level, we'll just use "u.u.u.1" to refer to the public IP address of the University and the IP address of the web server on the Internet will be "w.w.w.1".

Student A (IP address 10.1.1.56) sends an HTML request to a web server on the Internet with the IP address w.w.w.1. The applicable networking information in the packet breaks down as follows:

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1

Attribute	Original Packet	Packet after NATing
Destination IP address or dst-ip:	w . w . w . 1	w . w . w . 1
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

The source IP address is now that of the public facing interface of the FortiGate and source port number is an unused TCP port number on the FortiGate chosen by the FortiGate. Of these variable the only one the that FortiGate can really change and still have the packet reach the correct destination, in both directions, is the source port number.

There are a few methods of assigning the port number. First we'll look at the methods that are or have been used in the industry but aren't used by Fortinet.

Global pool

This method of differentiation focuses on the attribute of the source port number. In this approach a single pool of potential port numbers is set aside for the purposes of NAT. As a pool number is assigned, it is removed from the pool so that two sessions from different computers can not using the same port number. Once the session is over and no longer in use by the computer, the port number is put back into the pool where it can be assigned again.

Example global pool:

	Hexidecimal	Decimal
Start or range	0x7000	28672
End end of range	0xF000	61440
Possible ports in range	215	32768

This is a simple approach to implement and is good if the number of connections in unlike to reach the pool size. It would be okay for home use, but our example is for a university using 10 . 1 . 1 . 0 / 8 as a subnet. That means 16,777,214 possible IP addresses; more than this method can handle.

Fortinet does not use this method.

Global per protocol

This method uses the attributes source port number and type of protocol to differentiate between sessions. This approach is a variation of the first one. An additional piece of information is refered to in the packet that describes the protocol. For instance UDP or TCP. This could effectively double the number of potential addresses to NAT.

Example:

Here are two possible packets that would be considered different by the FortiGate so that any responses from the web server would make it back to their correct original sender.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	udp	udp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

Even though the source port is the same, because the protocol is different they are considered to be from different sessions and different computers.

The drawback is that it would depend on the protocols being used be evenly distributed between TCP and UDP. Even if this was the case the number would only double; reaching an upper limit of 65,536 possible connections. That number is still far short of the possible more than 16 million for an IP subnet with an eight bit subnet mask like the one in our example.

Fortinet does not use this method.

Per NAT IP Pool

This approach adds on to the previous one by adding another variable. In this case that variable is the IP addresses on the public side of the FortiGate. By having a pool of IP addresses to assign as the source IP address when NATing, the same number that was potentially available for the Global per protocol method can be multiplied by the number of external IP addresses in the pool. If you can assign a second IP address to the pool, you can double the potential number of sessions.

Example:

In this example it will be assumed that the FortiGate has 2 IP addresses that it can use. This could happen either by using two ISPs, or by having a pool of IP addresses assigned to a single interface. For simplicity will refer

to these IP public IP addresses as `u.u.u.1` and `u.u.u.2`.

Here are two possible packets that would be considered different by the FortiGate so that any responses from the web server would make it back to their correct original sender.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.2
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

In this example we even made the protocol the same. After the NATing process all of the variables are the same except the source address. This is still going to make it back to the original sender.

The drawback is that if you have only one IP address for the purposes of NATing this method does not gain you anything over the last method. Or if you do have multiple IP addresses to use it will still take quite a few to reach the 16 million possible that the subnet is capable of handling.

Fortinet does not use this method.

Per NAT IP, destination IP, port, and protocol

This is the approach that FortiOS uses.

It uses all of the differentiation point of the previous methods, NAT IP, port number and protocol, but the additional information point of the destination IP is also used. So now the network information points in the packet that the FortiGate keeps in its database to differentiate between sessions is:

- Public IP address of the FortiGate assigned by NATing
- Protocol of the traffic
- Source port assigned by the FortiGate
- Destination IP address of the packet

The last one is an especially good way to differentiate because as a theoretical number, the upper limit on that is the numbers of Public IP addresses on the whole of the Internet. Chances are that while a large number of sessions from inside the University will be going to a small group of sites such as Google, Youtube, Facebook and some others it is unlikely that they will all be going to them at the same time.

Example:

In this example it will be assumed that the FortiGate has only one IP address. Two possible packets will be described. The only difference in the attributes recorded will be the destination of the HTML request. These packets are still considered to be from different sessions and any responses will make it back to the correct computer.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.1
Destination IP address or dst-ip:	w.w.w.2	w.w.w.2
Protocol	tcp	tcp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

The reason that these attributes are used to determine differentiation between traffic is based on how the indexes for the sessions are recorded in the database. When a TCP connection is made through a FortiGate unit, a session is created and two indexes are created for the session. The FortiGate unit uses these indexes to guide matching traffic to the session.

This following could be the session record for the TCP connection in the first example.

Attribute	Outgoing Traffic	Returning Traffic
Source IP address	10.78.33.97 (internal address)	w.w.w.1
Destination address	w.w.w.1	u.u.u.1
Protocol	tcp	tcp
Source port	10000 (from original computer) 46372 (assigned by NAT)	80
Destination port	80	46372 (FortiGate assigned port)

Using the FortiGate's approach for session differentiation, FortiOS only has to ensure that the assigned port, along with the other four attributes is a unique combination to identify the session. So for example, if Student A simultaneously makes a HTTP(port 80) connection and a HTTPS(port 443) connection the same web server this would create another session and the index in the reply direction would be:

Attribute	Outgoing Traffic	Returning Traffic
Source IP address	10.78.33.97 (internal address)	w.w.w.1
Destination address	w.w.w.1	u.u.u.1
Protocol	tcp	tcp
Source port	10000 (from original computer) 46372 (assigned by NAT)	443
Destination port	443	46372 (FortiGate assigned port)

These two sessions are different and acceptable because of the different source port numbers on the returning traffic or the destination port depending on the direction of the traffic.

Calculations for possible session numbers

The result of using these four attributes instead of just the one that was originally used is a large increase in the number of possible unique combinations. For those who love math, the maximum number of simultaneous connections that can be supported is:

$$N \times R \times P \times D \times Dp$$

where:

- **N** is the number of NAT IP addresses
- **R** is the port range,

- **P** is the number of protocols,
- **D** is the number of unique destination IP addresses
- **Dp** the number of unique destination ports.

As a rough example let's do some basic calculations

- **N** - In our existing example we have already stated that there is only one public IP address that is being used by NAT. Realistically, for a university this number would likely be larger, but we're keeping it simple.

N = 1

R - The port range for our example has already been describe and we will keep it the same.

R = 32768

P - While there are a few protocols that are involved in Internet traffic we will limit this calculation just to TCP traffic.

P = 1

D - As mentioned before the number of unique destination addresses is growing larger every day, so figuring out the upper limit of that number would be difficult to say the least. Instead we will make the assumption that most of the university students, do to their shared interest and similar demographic will concentrate most of their web browsing to the same sites; sites such as YouTube, Facebook, Google, Twitter, Instagram, Wikipedia etc. This is not even taking into account the fact that many of these popular sites use load balancing and multiple IP addresses. As an arbitrary number let's use the number 25.

D = 25

Dp - To keep things simple it is tempting to limit the destination port to port 80, the one that many associate with web browsing, but this would not be realistic. the use of HTTPS, port 443 is on the rise. There is also email, DNS, FTP, NTP and a number of other background services that we use without thinking too closely about. Let's keep it small and say ten of them.

Dp = 10

The math on this very conservative calculation is:

$1 \times 32768 \times 1 \times 25 \times 10 = 8,192,000$ possible NAT sessions

When you take into account that the chances of everybody being online at the same time, going only to one of those 25 sites and not millions of others, and using only TCP not UDP or any of the other protocols, it starts to look like this method may provide enough potential unique sessions even for a subnet as large as the one described.

How packets are handled by FortiOS

To give you idea of what happens to a packet as it makes its way through the FortiGate unit here is a brief overview. This particular trip of the packet is starting on the Internet side of the FortiGate firewall and ends with the packet exiting to the Internal network. An outbound trip would be similar. At any point in the path if the packet is going through what would be considered a filtering process and if fails the filter check the packet is dropped and does not continue any further down the path.

This information is covered in more detail in other in the Troubleshooting chapter of the FortiOS Handbook in the Life of a Packet section.

The incoming packet arrives at the external interface. This process of entering the device is referred to as **ingress**.

Step #1 - Ingress

1. Denial of Service Sensor
2. IP integrity header checking
3. IPsec connection check
4. Destination NAT
5. Routing

Step #2 - Stateful inspection engine

1. Session Helpers
2. Management Traffic
3. SSL VPN
4. User Authentication
5. Traffic Shaping
6. Session Tracking
7. Policy lookup

Step #3 - Security profiles scanning process

1. Flow-based Inspection Engine
2. IPS
3. Application Control
4. Data Leak Prevention
5. Email Filter
6. Web Filter
7. Anti-virus
8. Proxy-based Inspection Engine
9. VoIP Inspection
10. Data Leak Prevention
11. Email Filter
12. Web Filter
13. Anti-virus
14. ICAP

Step #4 - Egress

1. IPsec
2. Source NAT
3. Routing

FortiGate modes

The FortiGate unit has a choice of modes that it can be used in, either NAT/Route mode or Transparent mode. The FortiGate unit is able to operate as a firewall in both modes, but some of its features are limited in Transparent mode. It is always best to choose which mode you are going to be using at the beginning of the set up. Once you start configuring the device, if you want to change the mode you are going to lose all configuration settings in the change process.

NAT/Route mode

NAT/Route mode is the most commonly used mode by a significant margin and is thus the default setting on the device. As the name implies the function of NAT is commonly used in this mode and is easily configured but there is no requirement to use NAT. The FortiGate unit performs network address translation before IP packets are sent to the destination network.

These are some of the characteristics of NAT/Route mode:

- Typically used when the FortiGate unit is a gateway between private and public networks.
- Can act as a router between multiple networks within a network infrastructure.
- When used, the FortiGate unit is visible to the networks that it is connected to.
- Each logical interface is on a distinct subnet.
- Each Interface needs to be assigned a valid IP address for the subnet that it is connected to it.

Transparent mode

Transparent mode is so named because the device is effectively transparent in that it does not appear on the network in the way that other network devices show as a nodes in the path of network traffic. Transparent mode is typically used to apply the FortiOS features such as Security Profiles etc. on a private network where the FortiGate unit will be behind an existing firewall or router.

These are some of the characteristics of Transparent mode:

- The FortiGate unit is invisible to the network.
- All of its interfaces are on the same subnet and share the same IP address.
- The FortiGate unit uses a Management IP address for the purposes of Administration.
- Still able to use NAT to a degree, but the configuration is less straightforward

In Transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools.

Quality of Service

The Quality of Service (QoS) feature allows the management of the level of service and preference given to the various types and sources of traffic going through the firewall so that the traffic that is important to the services and functions connecting through the firewall gets the treatment required to ensure the level of quality that is required. QoS can be helpful for organizations that are trying to manage their voice and streaming multi-media

traffic, which can rapidly consume bandwidth. Both voice and streaming multi-media are sensitive to latency. FortiGate units support QoS using traffic policing, traffic shaping, and queuing.

Traffic policing

Packets are dropped that do not conform to bandwidth limitations

Traffic shaping

Assigning minimum levels of bandwidth to be allocated to specific traffic flows to guarantee levels of service or assigning maximum levels of bandwidth to be allocated to specific traffic flows so that they do not impede other flows of traffic.

This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.

Queuing

Assigning differing levels of priority to different traffic flows so that traffic flows that are adversely effected by latency are prevented from being effected by traffic flows that are not subject to the effects of latency. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

An example of where you would want to use something like this is if you had competing traffic flows of Voice over IP traffic and email traffic. The VoIP traffic is highly susceptible to latency issues. If you have a delay of a few seconds it is quickly noticeable when it is occurring. Email on the other hand can have a time delay of much longer and it is highly unlikely that it will be noticed at all.



By default, the priority given to any traffic is high, so if you want to give one type of traffic priority over all other traffic you will need to lower the priority of all of the other traffic.

Interfaces and zones

A Firewall is a gateway device that may be the nexus point for more than 2 networks. The interface that the traffic is coming in on and should be going out on is a fundamental concern for the purposes of routing as well as security. Routing, policies and addresses are all associated with interfaces. The interface is essentially the connection point of a subnet to the FortiGate unit and once connected can be connected to other subnets.

Physical interfaces or not the only ones that need to be considered. There are also virtual interfaces that can be applied to security policies. VLANs are one such virtual interface. Interfaces if certain VPN tunnels are another.

Policies are the foundation of the traffic control in a firewall and the Interfaces and addressing is the foundation that policies are based upon. Using the identity of the interface that the traffic connects to the FortiGate unit tells the firewall the initial direction of the traffic. The direction of the traffic is one of the determining factors in deciding how the traffic should be dealt with. You can tell that interfaces are a fundamental part of the policies because, by default, this is the criteria that the policies are sorted by.

Zones are a mechanism that was created to help in the administration of the firewalls. If you have a FortiGate unit with a large number of ports and a large number of nodes in your network the chances are high that there is going to be some duplication of policies. Zones provide the option of logically grouping multiple virtual and physical FortiGate firewall interfaces. The zones can then be used to apply security policies to control the incoming and outgoing traffic on those interfaces. This helps to keep the administration of the firewall simple and maintain consistency.

For example you may have several floors of people and each of the port interfaces could go to a separate floor where it connects to a switch controlling a different subnet. The people may be on different subnets but in terms of security they have the same requirements. If there were 4 floors and 4 interfaces a separate policy would have to be written for each floor to be allowed out on to the Internet off the WAN1 interface. This is not too bad if that is all that is being done, but now start adding the use of more complicated policy scenarios with Security Profiles, then throw in a number of Identity based issues and then add the complication that people in that organization tend to move around in that building between floors with their notebook computers.

Each time a policy is created for each of those floors there is a chance of an inconsistency cropping up. Rather than make up an additional duplicate set of policies for each floor, a zone can be created that combines multiple interfaces. And then a single policy can be created that uses that zone as one side of the traffic connection.

Firewall objects

As was mentioned earlier, the components of the FortiGate firewall go together like interlocking building blocks. The Firewall objects are a prime example of those building blocks. They are something that can be configured once and then used over and over again to build what you need. They can assist in making the administration of the FortiGate unit easier and more intuitive as well as easier to change. By configuring these objects with their future use in mind as well as building in accurate descriptions the firewall will become almost self documenting. That way, months later when a situation changes, you can take a look at a policy that needs to change and use a different firewall object to adapt to the new situation rather than build everything new from the ground up to accommodate the change.

This chapter includes information about the following Firewall objects:

- Addresses
- Services and TCP ports
- Firewall schedules
- Security profiles

UUID support

A Universally Unique Identified (UUID) attribute has been added to some firewall objects, so that the logs can record these UUID to be used by a FortiManager or FortiAnalyzer unit. The objects currently include:

- Addresses, both IPv4 and IPv6
- Address Groups, both IPv4 and IPv6
- Virtual IPs, both IPv4 and IPv6
- Virtual IP groups, both IPv4 and IPv6
- Policies, IPv4, IPv6 and IP64

A UUID is a 16-octet (128-bit) number that is represented by 32 lowercase hexadecimal digits. The digits are displayed in five groups separated by hyphens (-). The pattern is 8-4-4-4-12; 36 digits if you include the hyphens.



Note: UUID is only supported on large-partition platforms ($\geq 128\text{M}$)

Addresses

Firewall addresses define sources and destinations of network traffic and are used when creating policies. When properly set up these firewall objects can be used with great flexibility to make the configuration of firewall policies simpler and more intuitive. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

The addresses in the FortiGate unit can include:

- IPv4 addresses
- IPv6 addresses

- IPv4 Address Groups
- IPv6 Address Groups
- IP Pools
- Virtual IP Addresses
- Geography based addresses
- Wildcard addresses and netmasks
- Fully Qualified Domain Name addresses
- Multicast addresses

Interfaces

When setting up an address one of the parameters that is asked for is the interface. This means that the system will expect to see that address only on the interface that you select. You can only select one interface. If you expect that the address may be seen at more than one interface you can choose the “any” interface option. Whenever, possible it is best to choose a more specific interface than the “any” option because in the GUI configuration of firewall policies there is a drop down field that will show the possible addresses that can be used. The drop down will only show those addresses that can be on the interface assigned for that interface in the policy.

Example:

- You have an address called “XYZ”.
- “XYZ” is set to the WAN1 interface because that is the only interface that will be able to access that address.
- When you are selecting a Source Address in the Web-based Manager for a policy that is using the DMZ the address “XYZ” will not be in the drop-down menu.

When there are only 10 or 20 addresses this is not a concern, but if there are a few hundred addresses configured it can make your life easier.

Addresses, address groups, and virtual IPs must have unique names to avoid confusion in firewall policies. If an address is selected in a policy, the address cannot be deleted until it is deselected from the policy.



Addressing Best Practices Tip

The other reason to assign a specific interface to addresses is that it will prevent you from accidentally assigning an address where it will not work properly. Using the example from earlier, if the “XYZ” address was assigned to the “Any” interface instead of WAN1 and you configure the “XYZ” address.

IPv4 addresses

When creating an IPv4 address there are a number of different types of addresses that can be specified. These include:

- Subnet
- IP Range
- FQDN
- Geography
- URL Pattern (Explicit Proxy)

Which one chosen will depend on which method most easily yet accurately describes the addresses that you are trying to include with as few entries as possible based on the information that you have. For instance, if you are trying to describe the addresses of a specific company's web server but if you have no idea of how extensive there webserver farm is you would be more likely to use a Fully Qualified Domain Name (FQDN) rather than a specific IP address. On the other hand some computers don't have FQDNs and a specific IP address must be used.

The following is a more comprehensive description of the different types of addresses.

Subnet addresses

The subnet type of address is expressed using a host address and a subnet mask. From a strictly mathematical stand point this is the most flexible of the types because it can refer to as little as the address for one computer or as many as all of them.

It is usually used when referring to your own internal addresses because you know what they are and they are usually administered in groups that are nicely differentiated along the lines of the old A, B, and C classes of IPv4 addresses. They are also addresses that are not likely to change with the changing of Internet Service Providers (ISP).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- A single host such as a single computer with the address 192.45.46.45
- A range of hosts such as all of the hosts on the subnet 192.45.46.1 to 192.45.46.255
- All hosts, represented by 0.0.0.0 which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- Netmask for a class A subnet of 16,777,214 usable addresses: 255.0.0.0, or /8
- Netmask for a class B subnet of 65,534 usable addresses: 255.255.0.0, or /16
- Netmask for a class C subnet of 254 usable addresses: 255.255.255.0, or /24
- Netmask for subnetted class C of 126 usable addresses: 255.255.255.128, or /25
- Netmask for subnetted class C of 62 usable addresses: 255.255.255.128, or /26
- Netmask for subnetted class C of 30 usable addresses: 255.255.255.128, or /27
- Netmask for subnetted class C of 14 usable addresses: 255.255.255.128, or /28
- Netmask for subnetted class C of 6 usable addresses: 255.255.255.128, or /29
- Netmask for subnetted class C of 2 usable addresses: 255.255.255.128, or /30
- Netmask for a single computer: 255.255.255.255, or /32
- Netmask used with 0.0.0.0 to include all IP addresses: 0.0.0.0, or /0

So for a single host or subnet the valid format of IP address and netmask could be either:

x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

or

x.x.x.x/x, such as 192.168.1.0/24

Creating a subnet address

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category**, “**Address**”.
4. Input a **Name** for the address object.
5. Select the **Type** **Subnet** from the dropdown menu.
6. Input the host address and subnet mask in the **Subnet / IP Range** field.
7. Select the **Interface** from the dropdown menu.
8. Check the **Visibility** box.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a Subnet address for a database server on the DMZ:

Category	Address
Name	DB_server_1
Type	Subnet
Subnet / IP Range	192.168.50.26
Interface	DMZ
Visibility	[checked]
Comments	Database server used by public facing company web servers

IP range addresses

Where the Subnet address is good at representing a standardized group of addresses that are subnets the IP Range type of address can describe a group of addresses while being specific and granular. It does this by specifying a continuous set of IP addresses between one specific IP address and another. While it is most common that this range is with a subnet it is not a requirement. For instance, 192.168.1.0/24 and 192.168.2.0/24 would be 2 separate subnets but if you wanted to describe the top half of one and the bottom half of the other you could describe the range of 192.168.1.128-192.168.2.127. It's also a lot easier that trying to calculate the correct subnet mask.

The format would be:

x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120

There is a notation that is commonly used and accepted by some devices that follows the format:

x.x.x.[x-x], such as 192.168.110.[100-120]

This format is not recognized in FortiOS 5.2 as a valid IP Range.

Creating a IP range address

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category**, “**Address**”.
4. Input a **Name** for the address object.
5. Select the **Type** **IP Range** from the dropdown menu.
6. Input the host address and subnet mask in the **Subnet / IP Range** field.
7. Select the **Interface** from the dropdown menu.
8. Check the **Visibility** box.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Category	Address
Name	Guest_users
Type	IP Range
Subnet / IP Range	192.168.100.200-192.168.100.240
Interface	Port1
Visibility	[checked]
Comments	Computers on the 1st floor used by guests for Internet access.



IP Range addresses can now be configured for IPv6 addresses. The only differences in creating an IPv6 IP Range address is that you would choose IPv6 Address for the Category and the syntax of the address in the Subnet/IP Range field would be in the format of 2001:0db8:0000:0002:0:0:0:20-2001:0db8:0000:0004:0:0:0:20

FQDN addresses

By using Fully Qualified Domain Name (FQDN) addressing you can take advantage of the dynamic ability of the service to keep up with address changes without having to manually change the addresses on the FortiGate. FQDN addresses are most often used with external web sites but they can be used for internal web sites as well if there is a trusted DNS server that can be accessed. FQDN addressing also comes in handy for large web sites

that may use multiple addresses and load balancers for their web sites. The FortiGate firewall automatically maintains a cached record of all the addresses resolved by the DNS for the FQDN addresses used.

For example, if you were doing this manually and you wanted to have a security policy that involved Google you could track down all of the IP addresses that they use across multiple countries. Using the FQDN address is simpler and more convenient.

When representing hosts by an FQDN, the domain name can also be a subdomain, such as mail.example.com.

Valid FQDN formats include:

- <host_name>.<top_level_domain_name> such as example.com
- <host_name>.<second_level_domain_name>.<top_level_domain_name>, such as mail.example.com

When creating FQDN entries it is important to remember that:

- Wildcards are not supported in FQDN address objects
- While there is a level of convention that would imply it, “www.example.com” is not necessarily the same address of “example.com”. they will each have their own records on the DNS server.

The FortiGate firewall keeps track of the DNS TTLs so as the entries change on the DNS servers the IP address will effectively be updated for the FortiGate. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache.



There is a possible security downside to using FQDN addresses. Using a fully qualified domain name in a security policy means that your policies are relying on the DNS server to be accurate and correct. DNS servers in the past were not seen as potential targets because the thinking was that there was little of value on them and therefore are often not as well protected as some other network resources. People are becoming more aware that the value of the DNS server is that in many ways it controls where users and computers go on the Internet. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category**, “**Address**”.
4. Input a **Name** for the address object.
5. Select the **TypeFQDN** from the dropdown menu.
6. Input the domain name in the **FQDN** field.
7. Select the **Interface** from the dropdown menu.
8. Check the **Visibility** box.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a FQDN address for a remote FTP server used by Accounting team:

Category	Address
Name	Payroll_FTP_server
Type	FQDN
FQDN	ftp.payrollcompany.com
Interface	any
Visibility	[checked]
Comments	Third party FTP server used by Payroll.

Geography based addresses

Geography based addressing enables you to design policies based on addresses that are associated with a country.

This feature can be used to make either inclusive or exclusive policies. For instance, if you have a SSL VPN where the users will only be connecting from a single country, but you don't know from where in that country, you can filter out any connections coming in that are from outside that country.

On the other side of the equation, if you find that you are constantly being attacked by malicious intruders from a few countries that you have no dealings with you can block access to them before any traffic comes through.

The matching of geographical country designations to an IP address is achieved by collecting data from any IP addresses that connect to any of the FortiGuard Servers throughout the world. As a secondary task, when a FortiGuard server connects to an IP it also does a search on the Country of origin for the address and updates the database.

There is no single comprehensive list of IP addresses and their locations available because IP addresses can be transferred between ISPs or countries and some organization may not keep complete or up-to-date records regarding locations. FortiGuard Services are constantly updating their database of addresses matched to locations, but the database is dynamic and there may be addresses that have not been resolved to a location. While this means that there can be gaps in the completeness of the database it is possible to fill them in manually by means local to your FortiGate unit.

FortiOS IPv6 does not support the creation of geography-based address objects. This feature is for IPv4 addresses only.

Best Practices Tip:



Based on the limitation of the IP address matched to country database, it is best to use this type of address in a group with other addresses to fill in the gaps. For instance, if you are a company in Country "A" and all of your employees that will be using the SSL-VPN connection are in that country, the best practice would be to create an address group that includes the geographical address of Country "A". As valid addresses appear that are not allowed, you can add these other IPs to that group using IP addresses or IP range addresses without having to change the policy itself.

If you are trying to block addresses the principle works just the same. Your logs show that someone from IP address x.x.x.x has been trying to connect inappropriately to your network. You use a IP locator web site to determine that they have been attempting to connect from Country "X". Up until now they have not been successful, but you don't deal with the country they are connecting from so don't mind blocking the whole country. Create an address group that is designed for Blocking Access to any addresses in it, then add the geographical address for Country "X". Even if the policy does not block every single IP address from that country you have greatly increased your odds of blocking potential intrusion attempts. As your logs show other attempts you can look them up in an IP locator web site and if they are from the same country you can add the IP address for the subnet that they are connecting from.

Creating a geography address

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category**, "**Address**".
4. Input a **Name** for the address object.
5. Select the **Type, Geography** from the dropdown menu.
6. Select the country from the dropdown menu in the **Country** field.

The bottom of the dropdown menu has a search field so that you don't have to scroll through the entire list of countries.

7. Select the **Interface** from the dropdown menu.
8. Check the **Visibility** box.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a Geography address for a country that should be able to access resources on the network.

Category	Address
Name	United States
Type	Geography
Country	United States
Interface	any
Visibility	[checked]
Comments	

URL pattern addresses

This address is different from the other addresses in that it is not designed to be used in the normal firewall policy configuration. It is intended to be used only with explicit web proxies, and then only as a destination. In some respects it is like a FQDN address in that it uses DNS to look up the IP address that represents but it then goes a little more granular and further filters by the URL that is used to specify a location within the website itself. Wildcard characters can even be used to keep the address from being too granular.

Creating a IP range address

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category**, “**Address**”.
4. Input a **Name** for the address object.
5. Select the **Type**, **URL Pattern (Explicit Proxy)** from the dropdown menu.
6. Enter the URL value.
7. Select the **Interface** from the dropdown menu.
8. Check the **Visibility** box.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a URL pattern address for a wiki on the example.com website.

Category	Address
Name	Example.com_wiki
Type	URL Pattern (Explicit Proxy)
URL Pattern	example.com/wiki/*
Interface	any
Visibility	[checked]
Comments	<optional>

Address groups

Address groups are designed for ease of use in the administration of the device. If you have a number of addresses or address ranges that will commonly be treated the same or require the same security policies, you can put them into address groups, rather than entering multiple individual addresses in each policy refers to them.

The use of groups is not required. If you have a number of different addresses you could add them individually to a policy and the FortiGate firewall will process them just as quickly and efficiently as if they were in a group, but the chances are that if you have used a group once you could need to use it again and depending on the number of addresses involved entering them individually for each policy can become tedious and the likelihood of an address being missed becomes greater. If you have a number of policies using that combination of addresses it is much easier to add or subtract addresses from the group than to try and remember all of the firewall policies that combination of addresses was used in. With the group, you only have to make the one edit and it is used by any firewall policy using that address group.

Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to Any.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of Any, they can be grouped, even if the addresses involve different networks.

IPv4 address groups and IPv6 address groups are created and treated separately. You cannot mix IPv4 firewall addresses and IPv6 firewall addresses in the same address group. Because the Internet is currently based on IPv4 addresses IPv6 address groups cannot include FQDN or Geography based addresses.

Creating an address group

The method for creating either a IPv4 group or an IPv6 group is identical except for the selection of **Type**.

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select the down arrow next to Create New, select **Address Group**.
3. Choose the **Type**, that is applicable to the proposed selection of addresses.
4. Input a **Group Name** for the address object.
5. Check the **Show in Address List** box.
6. Next to **Members** there is a dropdown menu that can be used to select from the available Address objects. It is possible to select more than 1 entry. Just select the green plus sign next to the field to add an additional entry. Select the "X" icon in the field to remove an entry.
7. Input any additional information in the **Comments** field.
8. Press **OK**.

UUID support

Syntax:

```
config firewall {address|addres6|addgrp|addgrp6}
edit 1
set uuid <example uuid: 8289ef80-f879-51e2-20dd-fa62c5c51f44>
```

```
    next
end
```

Virtual IPs

The mapping of a specific IP address to another specific IP address is usually referred to as Destination NAT. FortiOS has a component that is a bit more specialized along this line called a Virtual IP Address, sometimes referred to as a VIP. FortiOS uses a Virtual IP address to map an External IP address to an IP address. This address does not have to be an individual host, it can also be an address range. This mapping can include all TCP/UDP ports or if Port Forwarding is enabled it will only refer to the specific ports configured.

Virtual IP addresses are typically used to NAT external or Public IP addresses to internal or Private IP addresses. Using a Virtual IP address between 2 internal Interfaces made up of Private IP addresses is possible but there is rarely a reason to do so as the 2 networks can just use the IP addresses of the networks without the need for any address translation. Using a Virtual IP address for traffic going from the inside to the Internet is even less likely to be a requirement, but it is supported.

Something that needs to be considered when there are multiple Public IP addresses on the external interface(s) is that when a Virtual IP address is used without Port Forwarding enabled there is a reciprocal effect as far as traffic flow is concerned. Normally, on a firewall policy where NAT is enabled, for outgoing traffic the internal address is translated to the Public address that is assigned to the FortiGate, but if there is a Virtual IP address with no port forwarding enabled, then the Internal IP address in the Mapped field would be translated to the IP address configured as the External Address in the VIP settings.



Best practice: Put any policies with a VIP right at the beginning of the policy list, with nothing before them. VIP traffic is processed first, before the regular rules of order in policies are applied. The sequence of the policies not containing VIPs will not effect those that do contain VIPs, regardless of the order they are in. Put VIP policies before any others to remind yourself of where they really are in the sequence. For more on this topic, read the ["Exception to policy order" on page 93](#).

Example

- The assigned External address (WAN1) of the FortiGate unit is 172.12.96.3 with a subnet mask of 255.255.255.128
- There is a Virtual IP address set up to map the external address 172.12.96.127 on WAN1 to the internal IP address of 192.168.1.127
- Port Forwarding is not enabled because you want all allowed traffic going to the external IP address to go to this server.

In this case any outbound traffic from 192.168.1.127 will go out on WAN1 with the IP address of 172.12.96.127 as the source IP address.

In terms of actually using the Virtual IP address, they would be using in the security policies in the same places that other addresses would be used, usually as a Destination Address.

UUID support for VIP

UUID is now supported in for virtual IPs and virtual IP groups. This includes virtual IPs for IPv4, IPv6, NAT46, and NAT64. To view the UUID for these objects in a FortiGate unit's logs, log-uuid must be set to extended mode, rather than policy-only (which only shows the policy UUID in a traffic log). UUID can only be configured through the CLI

Syntax:


```
config sys global
  set log-uuid {disable | policy-only | extended}
end
```



There is another type of address that the term “virtual IP address” commonly refers to which is used in load balancing and other similar configurations. In those cases, a number of devices share a separately created virtual IP address that can be sent to multiple possible devices. In FortiOS these are referred to as Virtual Servers and are configured in the “Load Balance” section.

Dynamic VIP according to DNS translation

When a dynamic virtual IP is used in a policy, the dynamic DNS translation table is installed along with the dynamic NAT translation table into the kernel. All matched DNS responses will be translated and recorded regardless if they hit the policy. When a client request hits the policy, dynamic NAT translation will occur if it matches a record, otherwise the traffic will be blocked.

Syntax

```
config firewall vip
  edit "1"
    set type dns-translation
    set extip 192.168.0.1-192.168.0.100
    set extintf "dmz"
    set dns-mapping-ttl 604800
    set mappedip "3.3.3.0/24" "4.0.0.0/24"
  end
end
```

Creating a virtual IP

1. Go to **Policy & Objects > Objects > Virtual IPs**.
2. Select **Create New**.
 - If you use the down arrow next to Create New, select **Virtual IP**.
3. Choose the **VIP Type**.

The options available are:

- a. IPv4 VIP - IPv4 on both sides of the FortiGate Unit.
- b. IPv6 VIP - IPv6 on both sides of the FortiGate Unit.
- c. NAT46 VIP - Going from an IPv4 Network to an IPv6 Network.
- d. NAT64 VIP - Going from an IPv6 Network to an IPv4 Network.

Which is chosen will depend on which of the IP version networks is on the external interface of the FortiGate unit and which is on the internal interface.

4. Input a **Name** for the Virtual IP.
5. Input any additional information in the **Comments** field.
6. Using the dropdown menu for the Interface Field, choose the incoming interface for the traffic.

The IPv4 VIP Type is the only one that has a field for the interface. This is a legacy function from previous versions so that they can be upgraded without complicated reconfiguration. The External IP address, which is a required

field, tells the unit which interface to use so it is perfectly acceptable to choose **Any** as the interface. In some configurations, if the Interface field is not set to **Any** the VIP is not one of the displayed options when choosing a destination address.

7. If only specific IP addresses are allowed to be the source address for traffic using the VIP, check the box for the **Source Address Filter**.
 - a. To specify the allowed address range select **Create New**.
 - b. Enter the IP address for the start of the set of IP address(es) in the **Range Start** field.
 - c. Enter the IP address for the end of the set of IP address(es) in the **Range End** field.
8. Enter the IP address for the **External IP Address/Range**.
If there is a single IP address, use that address in both fields.
9. Set the **Mapped IP Type**.
This will be either Subnet or Address Range.
If you only have a single destination address you can use either:
 - **Subnet:** x.x.x.x/32
 - **Address Range:** x.x.x.x - x.x.x.x, where x.x.x.x is the same IP address.
10. Enter the IP address(es) for the **Mapped IP Address/Range**.
This will be the address of the host that the traffic is being directed to.
11. If you are only going to use specific ports, enable **Port Forwarding**.
 - a. Select one of 3 **Protocol** types:
 - **TCP**
 - **UDP**
 - **SCTP**
 - b. Enter the port number or range that the traffic will be connecting to in the **External Service Port** fields.
 - c. Enter the port number or range that is the final destination of the traffic in the **Map to Port** fields
12. Press **OK**.

Example

This example is for a VIP that is being used to direct traffic from the external IP address to a webserver on the internal network. The webserver is for company use only. The company's public facing webserver already used port 80 and there is only one IP external IP address so the traffic for this server is being listened for on port 8080 of the external interface and being sent to port 80 on the internal host.

VIP Type	IPv4
Name	Internal_Webserver
Comments	Webserver with Collaboration tools for Corporate employees
Interface	Any
Source Address Filter	<list of IP addresses of remote users>

External IP Address/Range	172.13.100.27 <this would normally be a public IP address
Mapped IP Type	Subnet
Mapped IP Address/Range	192.168.34.150
Port Forwarding	enabled
Protocol	TCP
External Service Port	8080 - 8080
Map to Port	80 - 80

Virtual IP groups

Just like other address, Virtual IP addresses can be organized into groups for ease of administration. If you have multiple virtual IPs that are likely to be associated to common firewall policies rather than add them individually to each of the policies you can add the instead. That way, if the members of the group change then any changes made to the group will propagate to all of the policies using that group.

When using a Virtual IP address group the firewall policy will take into account all of the configured parameters of the Virtual IPs: IP addresses, Ports and port types.

Creating a virtual IP group

1. Go to **Policy & Objects > Objects > Virtual IPs**.
2. Use the down arrow next to Create New, select **Virtual IP Group**.
3. Select the **Type** fo VIP group you wish to create.

The options available are:

- a. **IPv4 VIP** - IPv4 on both sides of the FortiGate Unit.
- b. **IPv6 VIP** - IPv6 on both sides of the FortiGate Unit.
- c. **NAT46 VIP** - Going from an IPv4 Network to an IPv6 Network.
- d. **NAT64 VIP** - Going from an IPv6 Network to an IPv4 Network.

Which is chosen will depend on which of the IP version networks is on the external interface of the FortiGate unit and which is on the internal interface. The options will be:

4. Enter a **Name** for the VIP
5. Enter any additional information in the **Comments** field.
6. Use the dropdown menu of the **Interface** field to select the interface that the member VIPs are associated with, if any.
7. Use the dropdown menu fo the **Members** field to select the VIPs that will be part of the group.
8. Press **OK**.

IP pools

IP pools are a mechanism that allow sessions leaving the FortiGate Firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses will be used instead of the IP address assigned to that FortiGate interface.

There are 4 types of IP Pools that can be configured on the FortiGate firewall:

- One-to-One - in this case the only internal address used by the external address is the internal address that it is mapped to.
- Overload - this is the default setting. Internal addresses other than the one designated in the policy can use this address for the purposes of NAT.
- Fixed Port Range - rather than a single address to be used, there is a range of addresses that can be used as the NAT address. These addresses are randomly assigned as the connections are made.
- Port Block Allocation - this setting is used to allocate a block of port numbers for IP pool users. Two variables will also have to be set. The block size can be set from 64 to 4096 and as the name implies describes the number of ports in one block of port numbers. The number of blocks per user determines how many of these blocks will be assigned. This number can range from 1 to 128.



Be careful when calculating the values of the variables. The maximum number of ports that are available on an address is 65,536. If you chose the maximum value for both variables you will get a number far in excess of the available port numbers.

$$4096 \times 128 = 524,288$$

One of the more common examples is when you have an email server behind your FortiGate firewall and the range of IP addresses assigned to you by your ISP is more than one. If an organization is assigned multiple IP addresses it is normally considered a best practice to assign a specific address other than the one used for the Firewall to the mail server. However, when normal NAT is used the address assigned to the firewall is also assigned to any outbound sessions. Anti-spam services match the source IP address of mail traffic that they receive to the MX record on DNS servers as an indicator for spam. If there is a mismatch the mail may not get through so there is a need to make sure that the NATed address assigned matches the MX record.

You can also use the Central NAT table as a way to configure IP pools.

Source IP address and IP pool address matching when using a range

When the source addresses are translated to an IP pool that is a range of addresses, one of the following three cases may occur:

Scenario 1:

The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable fixed port in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Scenario 2:

The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism. If you enable fixed port in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Scenario 3:

The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

ARP replies

If a FortiGate firewall interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools. For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

$(1.1.1.0-1.1.1.255) \text{ and } (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20$

The port2 interface overlap IP range with IP_pool_2 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20$

The port2 interface overlap IP range with IP_pool_3 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40$

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select Enable NAT in a security policy and then select Dynamic IP Pool. Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool. Whether or not the external address of an IP Pool will respond to an ARP request can be disabled. You might want to disable the ability to responded to ARP requests so that these address cannot be used as a way into your network or show up on a port scan.

IP pools and zones

Because IP pools are associated with individual interfaces

IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

Creating a dynamic IP pool

A dynamic IP pool is essentially one in which the IP address that is assigned to the sending computer is not known until the session is created, therefore at the very least it will have to be a pool of at least 2 potential addresses. A quick example would be an IP pool for users of a VPN.

IP Pool Type	IPv4 Pool
Name	Sales_Group_VPN_Users
Comments	<optional>
Type	One-to-One
External IP Range	10.23.0.50 - 10.23.0.75
ARP Reply	[checked]

Fixed port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

However, enabling the use of a fixed port means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Match-VIP

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. By default, the feature is disabled.

Services and TCP ports

There are a number of different services and protocols in use on the Internet. The most commonly known is HTTP which is used by web servers to transmit requests and responses for unencrypted web pages. These services are set up to listen for requests on a numbered port. These services and protocols can use any port from 1 to 65,535. To keep things simple for everyone a large number of the more commonly used services started using a standardized list of ports. For instance, though it is not required, by default, most web servers listen for HTTP requests on port 80 and by default, web browsers will send HTTP traffic to port 80. If you wish to use another port such as 8080 you would put “:8080” at the end of the URL to indicate that you want the browser to use 8080 instead of the default port.

Example

Default URL for HTTP traffic when the web server is listening on the standard HTTP port:

`http://fortinet.com`

URL to the same address when the web server is listening for HTTP traffic on port 8080

`http://fortinet.com:8080`

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined on the FortiGate unit. If there is a service that does not appear on the list you can create a service or edit an existing one. You need to know the ports, IP addresses or protocols of that particular service or application uses, to create a service.

Best Practices



While you can edit a predefined service it is best to leave those ones alone and create a new service and name it something similar such as the same service name with a descriptive identifier appended.

Based on the previous example, instead of the name “HTTP” you could name the service “HTTP8080” or use the application that is using that port, “HTTP-Application”.

Categories

In order to make sorting through the services easier there is a field to categorize the services. The services can be sorted into the following groups:

- Uncategorized
- General
- Web Access
- File Access
- Email

- Network Services
- Authentication
- Remote Access
- Tunnelling
- VoIP, Messaging and Other Applications
- Web Proxy

To create a new category, use the downward pointing arrow next to **Create New** in the **Services** window and choose **Category**. All that will be required is a name for the new category. A comments describing the new category is optional.

Protocol types

One of the fundamental aspects of a service is the type of protocol that use used to define it. When a service is defined one of the following categories of protocol needs to be determined:

- TCP/UDP/SCTP
- ICMP
- ICMP6
- IP

Depending on which of these protocol categories is choose another set of specifications will can also be defined.

TCP/UDP/SCTP

This is the most commonly used service protocol category. Once this category has been selected the other available options to choose are an address, either IP or FQDN, and the protocol and port number.

The protocol will be TCP, UDP or SCTP.

ICMP or ICMP6

When ICMP or ICMP6 is chosen the available options are the ICMP Type and its code.

IP

When IP is the chosen protocol type the addition option is the Protocol Number.

TCP

Transmission Control Protocol (TCP) is one of the core or fundamental protocols of the Internet. It is part of the Transport Layer of the OSI Model. It is designed to provide reliable delivery of data from a program on one device on the network or Internet to another program on another device on the network or Internet. TCP achieves its reliability because it is a connection based protocol. TCP is stream-oriented. It transports streams of data reliably and in order.

TCP establishes a prior connection link between the hosts before sending data. This is often referred to as the handshake. Once the link is established the protocol uses checks to verify that the data transmitted. If an error check fails the data is retransmitted. This makes sure that the data is getting to the destination error free and in the correct order so that it can be put back together into a form that is identical to the way they were sent.

TCP is configured more for reliability than for speed and because of this TCP will likely be slower than a connectionless protocol such as UDP. This is why TCP is generally not used for real time applications such as voice communication or online gaming.

Some of the applications that use TCP are:

- World Wide Web (HTTP and HTTPS)
- Email (SMTP, POP3, IMAP4)
- Remote administration (RDP)
- File transfer (FTP)

UDP

User Datagram Protocol (UDP) like TCP is one of the core protocols of the Internet and part of the Transport Layer of the OSI Model. UDP is designed more for speed than reliability and is generally used for different applications than TCP. UDP sends messages, referred to as datagrams across the network or Internet to other hosts without establishing a prior communication link. In other words, there is no handshake.

UDP is an unreliable service as the datagrams can arrive out of order, duplicated or go missing without any mechanism to verify them. UDP works on the assumption that any error checking is done by the application or is not necessary for the function of the application. This way it avoids the overhead that is required to verify the integrity of the data.

This lack of overhead improves the speed of the data transfer and is why UDP is often used by applications that are time sensitive in nature. UDP's stateless nature is also great for applications that answer a large number of small queries from a large number of clients.

Common uses for UDP are:

- Domain Name Resolution (DNS)
- Time (NTP)
- Streaming media (RTSP, RTP and RTCP)
- Telephone of the Internet (VoIP)
- File Transfer (TFTP)
- Logging (SNMP)
- Online games (GTP and OGP)

SCTP

Stream Control Transmission Protocol (SCTP) is part of the Transport Layer of the OSI Model just like TCP and UDP and provides some of the features of both of those protocols. It is message or datagram orientated like UDP but it also ensures reliable sequential transport of data with congestion control like TCP.

SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- Network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP uses multi-streaming to transport its messages which means that there can be several independent streams of messages traveling in parallel between the points of the transmission. The data is sent out in larger chunks of data than is used by TCP just like UDP but the messages include a sequence number within each message in the same way that TCP does so that the data can be reassembled at the other end of the transmission in the correct sequence without the data having to arrive in the correct sequence.

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a much newer protocol. It was defined by the IETF Signaling Transport (SIGTRAN) working group in 2000. It was introduced by RFC 3286 and more fully defined by RFC 4960.

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to "ALL". FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism
- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPsec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

Specific addresses in TCP/UDP/SCTP

In the TCP/UDP/SCTP services it is also possible to set the parameter for a specific IP or Fully Qualified Domain Name address. The IP/FQDN field refers to the destination address of the traffic, not the source. This means for example, that you can set up a custom service that will describe in a policy the TCP traffic over port 80 going to

the web site example.com, but you cannot set up a service that describes the TCP traffic over port 80 that is coming from the computer with the address 192.168.29.59.

Protocol port values

The source and destination ports for TCP/UDP/SCTP services are important to get correct. If they are reversed the service will not work. The destination port(s) are the ones that refer to the ports that the computer will be listening on. These are the port numbers that most people are familiar with when they associate a port number to a protocol. In most cases the source port will be one that is randomly assigned by the computer that is not being already used by another service.

Most people associate HTTP with port 80. This means that a web-server will be listening on port 80 for any http requests being sent to the computer. The computer that is sending the request can use any port that is not already assigned to another service or communication session. There are 65,535 ports that it can randomly assign, but because the ports from 1 to 1024 are normally used for listening for incoming communications it is usually not in that range. It is unless there is a specific instance when you know that a communication will be coming from a predefined source port it is best practice to set the source port range from 1 to 65,535.

ICMP

The Internet Control Message Protocol (ICMP) is a protocol layered onto the Internet Protocol Suite to provide error reporting flow control and first-hop gateway redirection. It is normally used by the operating systems of networked computers to send connectivity status query, response and error messages. It is assigned protocol number 1. There is a version of the protocol for both IPv4 and for IPv6. It is not designed to be absolutely reliable like TCP.

ICMP is not typically used for transporting data or for end-user network applications with the exception of some diagnostic utilities such as ping and traceroute.

ICMP messages are sent in several situations, for example:

- when a datagram cannot reach its destination,
- time exceeded messages
- redirect messages
- when the gateway does not have the buffering capacity to forward a datagram
- when the gateway can direct the host to send traffic on a shorter route.

Some of the specific ICMP message types are:

- ICMP_ECHO
- ICMP_TIMESTAMP
- ICMP_INFO_REQUEST
- ICMP_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

ICMP types and codes

ICMP has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

ICMP types and codes

Type Number	Type Name	Optional Code(s)
0	Echo Reply	
1	Unassigned	
2	Unassigned	
3	Destination Unreachable	0 Net Unreachable
		1 Host Unreachable
		2 Protocol Unreachable
		3 Port Unreachable
		4 Fragmentation Needed and Don't Fragment was Set
		5 Source Route Failed
		6 Destination Network Unknown
		7 Destination Host Unknown
		8 Source Host Isolated
		9 Communication with Destination Network is Administratively Prohibited
		10 Communication with Destination Host is Administratively Prohibited
		11 Destination Network Unreachable for Type of Service
		12 Destination Host Unreachable for Type of Service
		13 Communication Administratively Prohibited
		14 Host Precedence Violation
		15 Precedence cutoff in effect

Type Number	Type Name	Optional Code(s)
4	Source Quench	
5	Redirect	0 Redirect Datagram for the Network (or subnet) 1 Redirect Datagram for the Host 2 Redirect Datagram for the Type of Service and Network 3 Redirect Datagram for the Type of Service and Host
6	Alternate Host Address	
7	Unassigned	
8	Echo	
9	Router Advertisement	
10	Router Selection	
11	Time Exceeded	0 Time to Live exceeded in Transit 1 Fragment Reassembly Time Exceeded
12	Parameter Problem	0 Pointer indicates the error 1 Missing a Required Option 2 Bad Length
13	Timestamp	
14	Timestamp Reply	
15	Information Request	
16	Information Reply	
17	Address Mask Request	
18	Address Mask Reply	

Type Number	Type Name	Optional Code(s)
19	Reserved (for Security)	
20 - 29	Reserved (for Robustness Experiment)	
30	Traceroute	
31	Datagram Conversion Error	
32	Mobile Host Redirect	
33	IPv6 Where-Are-You	
34	IPv6 I-Am-Here	
35	Mobile Registration	
36	Mobile Registration Reply	
37	Domain Name Request	
38	Domain Name Reply	
39	SKIP	
40	Photuris	
41 - 255	Reserved	

log-invalid-packet

The `log-invalid-packet` CLI setting is one that is intended to log invalid ICMP packets. The exact definition being:

If the FortiGate unit receives an ICMP error packet that contains an embedded `IP (A, B) | TCP (C, D)` header, then if FortiOS can locate the A:C -> B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped.

When this field is enabled, the FortiGate also log messages that are not ICMP error packets.

Types of logs covered by `log-invalid-packet`

- Invalid ICMP
 - If ICMP error message verification (see "check-reset-range") is enabled
- Invalid DNS packets
 - DNS packets that contain requests for non-existing domains
- `iprope` check failed
- reverse path check fail
- denied and broadcast traffic
- no session matched

Some other examples of messages that are not errors that will be logged, based on [RFC792](#):

Type 3 messages correspond to "Destination Unreachable Message"

- Type 3, Code 1 = host unreachable
- Type 3, Code 3 = port unreachable

Type 11 messages correspond to "Time Exceeded Message"

- Type 11, Code 0 = time to live exceeded in transit

ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) is the new implementation of the Internet Control Message Protocol (ICMP) that is part of Internet Protocol version 6 (IPv6). The ICMPv6 protocol is defined in [RFC 4443](#).

ICMPv6 is a multipurpose protocol. It performs such things as:

- error reporting in packet processing
- diagnostic functions
- Neighbor Discovery process
- IPv6 multicast membership reporting

It also designed as a framework to use extensions for use with future implementations and changes.

Examples of extensions that have already been written for ICMPv6:

- Neighbor Discovery Protocol (NDP) - a node discovery protocol in IPv6 which replaces and enhances functions of ARP.
- Secure Neighbor Discovery Protocol (SEND) - an extension of NDP with extra security.
- Multicast Router Discovery (MRD) - allows discovery of multicast routers.

ICMPv6 messages use IPv6 packets for transportation and can include IPv6 extension headers. ICMPv6 includes some of the functionality that in IPv4 was distributed among protocols such as ICMPv4, ARP (Address Resolution Protocol), and IGMP (Internet Group Membership Protocol version 3).

ICMPv6 has simplified the communication process by eliminating obsolete messages.

ICMPv6 messages are subdivided into two classes: error messages and information messages.

Error Messages are divided into four categories:

1. Destination Unreachable
2. Time Exceeded

3. Packet Too Big
4. Parameter Problems

Information messages are divided into three groups:

1. Diagnostic messages
2. Neighbor Discovery messages
3. Messages for the management of multicast groups.

ICMPv6 types and codes

ICMPv6 has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

Type codes 0 – 127 are error messages and type codes 128 – 255 are for information messages.

ICMPv6 Types and Codes

Type Number	Type Name	Code
0	Reserved	0 - no route to destination
		1 - communication with destination administratively prohibited
		2 - beyond scope of source address
		3 - address unreachable
		4 - port unreachable
		5 - source address failed ingress/egress policy
		6 - reject route to destination
		7 - Error in Source Routing Header
1	Destination Unreachable	
2	Packet Too Big	
3	Time Exceeded	0 - hop limit exceeded in transit
		1 - fragment reassembly time exceeded
4	Parameter Problem	0 - erroneous header field encountered
		1 - unrecognized Next Header type encountered
		2 - unrecognized IPv6 option encountered

Type Number	Type Name	Code
100	Private Experimentation	
101	Private Experimentation	
102 - 126	Unassigned	
127	Reserved for expansion if ICMPv6 error messages	
128	Echo Request	
129	Echo Replay	
130	Multicast Listener Query	
131	Multicast Listener Report	
132	Multicast Listener Done	
133	Router Solicitation	
134	Router Advertisement	
135	Neighbor Solicitation	
136	Neighbor Advertisement	
137	Redirect Message	
138	Router Renumbering	0 - Router Renumbering Command
		1 - Router Renumbering Result
		255 - Sequence Number Reset

Type Number	Type Name	Code
139	ICMP Node Information Query	0 - The Data field contains an IPv6 address which is the Subject of this Query.
		1 - The Data field contains a name which is the Subject of this Query, or is empty, as in the case of a NOOP.
		2 - The Data field contains an IPv4 address which is the Subject of this Query.
140	ICMP Node Information Response	0 - A successful reply. The Reply Data field may or may not be empty.
		1 - The Responder refuses to supply the answer. The Reply Data field will be empty.
		2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.
141	Inverse Neighbor Discovery Solicitation Message	
142	Inverse Neighbor Discovery Advertisement Message	
143	Version 2 Multicast Listener Report	
144	Home Agent Address Discovery Request Message	
145	Home Agent Address Discovery Reply Message	
146	Mobile Prefix Solicitation	
147	Mobile Prefix Advertisement	
148	Certification Path Solicitation Message	

Type Number	Type Name	Code
149	Certification Path Advertisement Message	
150	ICMP messages utilized by experimental mobility protocols such as Seamoby	
151	Multicast Router Advertisement	
152	Multicast Router Solicitation	
153	Multicast Router Termination	
154	FMIPv6 Messages	
155	RPL Control Message	
156	ILNIPv6 Locator Update Message	
157	Duplicate Address Request	
158	Duplicate Address Confirmation	
159 – 199	Unassigned	
200	Private experimentation	
201	Private experimentation	
255	Reserved for expansion of ICMPv6 informational messages	

IP

Internet Protocol (IP) is the primary part of the Network Layer of the OSI Model that is responsible for routing traffic across network boundaries. It is the protocol that is responsible for addressing. IPv4 is probable the version that most people are familiar with and it has been around since 1974. IPv6 is its current successor and due to a shortage of available IPv4 addresses compared to the explosive increase in the number of devices that use IP addresses, IPv6 is rapidly increasing in use.

When IP is chosen as the protocol type the available option to further specify the protocol is the protocol number. This is used to narrow down which protocol within the Internet Protocol Suite and provide a more granular control.

Protocol number

IP is responsible for more than the address that it is most commonly associated with and there are a number of associated protocols that make up the Network Layer. While there are not 256 of them, the field that identifies them is a numeric value between 0 and 256.

In the Internet Protocol version 4 (IPv4) [RFC791] there is a field called "Protocol" to identify the next level protocol. This is an 8 bit field. In Internet Protocol version 6 (IPv6) [RFC2460], this field is called the "Next Header" field.

Protocol Numbers

#	Protocol	Protocol's Full Name
0	HOPOPT	IPv6 Hop-by-Hop Option
1	ICMP	Internet Control Message Protocol
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway
4	IPv4	IPv4 encapsulation Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
7	CBT	CBT
8	EGP	Exterior Gateway Protocol
9	IGP	Any private interior gateway (used by Cisco for their IGRP)
10	BBN-RCC-MON	BBN RCC Monitoring
11	NVP-II	Network Voice Protocol
12	PUP	PUP

#	Protocol	Protocol's Full Name
13	ARGUS	ARGUS
14	EMCON	EMCON
15	XNET	Cross Net Debugger
16	CHAOS	Chaos
17	UDP	User Datagram Protocol
18	MUX	Multiplexing
19	DCN-MEAS	DCN Measurement Subsystems
20	HMP	Host Monitoring
21	PRM	Packet Radio Measurement
22	XNS-IDP	XEROX NS IDP
23	TRUNK-1	Trunk-1
24	TRUNK-2	Trunk-2
25	LEAF-1	Leaf-1
26	LEAF-2	Leaf-2
27	RDP	Reliable Data Protocol
28	IRTP	Internet Reliable Transaction
29	ISO-TP4	ISO Transport Protocol Class 4
30	NETBLT	Bulk Data Transfer Protocol
31	MFE-NSP	MFE Network Services Protocol
32	MERIT-INP	MERIT Internodal Protocol
33	DCCP	Datagram Congestion Control Protocol
34	3PC	Third Party Connect Protocol
35	IDPR	Inter-Domain Policy Routing Protocol
36	XTP	XTP

#	Protocol	Protocol's Full Name
37	DDP	Datagram Delivery Protocol
38	IDPR-CMTP	IDPR Control Message Transport Proto
39	TP++	TP++ Transport Protocol
40	IL	IL Transport Protocol
41	IPv6	IPv6 encapsulation
42	IPv6	SDRPSource Demand Routing Protocol
43	IPv6-Route	Routing Header for IPv6
44	IPv6-Frag	Fragment Header for IPv6
45	IDRP	Inter-Domain Routing Protocol
46	RSVP	Reservation Protocol
47	GRE	General Routing Encapsulation
48	DSR	Dynamic Source Routing Protocol
49	BNA	BNA
50	ESP	Encap Security Payload
51	AH	Authentication Header
52	I-NLSP	Integrated Net Layer Security TUBA
53	SWIPE	IP with Encryption
54	NARP	NBMA Address Resolution Protocol
55	MOBILE	IP Mobility
56	TLSP	Transport Layer Security Protocol using Kryptonet key management
57	SKIP	SKIP
58	IPv6-ICMP	ICMP for IPv6
59	IPv6-NoNxt	No Next Header for IPv6
60	IPv6-Opts	Destination Options for IPv6

#	Protocol	Protocol's Full Name
61		any host internal protocol
62	CFTP	CFTP
63		any local network
64	SAT-EXPAK	SATNET and Backroom EXPAK
65	KRYPTOLAN	Kryptolan
66	RVD	MIT Remote Virtual Disk Protocol
67	IPPC	Internet Pluribus Packet Core
68		any distributed file system
69	SAT-MON	SATNET Monitoring
70	VISA	VISA Protocol
71	IPCV	Internet Packet Core Utility
72	CPNX	Computer Protocol Network Executive
73	CPHB	Computer Protocol Heart Beat
74	WSN	Wang Span Network
75	PVP	Packet Video Protocol
76	BR-SAT-MON	Backroom SATNET Monitoring
77	SUN-ND	SUN ND PROTOCOL-Temporary
78	WB-MON	WIDEBAND Monitoring
79	WB-EXPAK	WIDEBAND EXPAK
80	ISO-IP	ISO Internet Protocol
81	VMTP	VMTP
82	SECURE-VMTP	SECURE-VMTP
83	VINES	VINES
84	TTP	TTP

#	Protocol	Protocol's Full Name
84	IPTM	Protocol Internet Protocol Traffic
85	NSFNET-IGP	NSFNET-IGP
86	DGP	Dissimilar Gateway Protocol
87	TCF	TCF
88	EIGRP	EIGRP
89	OSPFIGP	OSPFIGP
90	Sprite-RPC	Sprite RPC Protocol
91	LARP	Locus Address Resolution Protocol
92	MTP	Multicast Transport Protocol
93	AX.25	AX.25 Frames
94	IPIP	IP-within-IP Encapsulation Protocol
95	MICP	Mobile Internetworking Control Pro.
96	SCC-SP	Semaphore Communications Sec. Pro.
97	ETHERIP	Ethernet-within-IP Encapsulation
98	ENCAP	Encapsulation Header
99		any private encryption scheme
100	GMTP	GMTP
101	IFMP	Ipsilon Flow Management Protocol
102	PNNI	PNNI over IP
103	PIM	Protocol Independent Multicast
104	ARIS	ARIS
105	SCPS	SCPS
106	QNX	QNX
107	A/N	Active Networks

#	Protocol	Protocol's Full Name
108	IPComp	IP Payload Compression Protocol
109	SNP	Sitara Networks Protocol
110	Compaq-Peer	Compaq Peer Protocol
111	IPX-in-IP	IPX in IP
112	VRRP	Virtual Router Redundancy Protocol
113	PGM	PGM Reliable Transport Protocol
114		any 0-hop protocol
115	L2TP	Layer Two Tunneling Protocol
116	DDX	D-II Data Exchange (DDX)
117	IATP	Interactive Agent Transfer Protocol
118	STP	Schedule Transfer Protocol
119	SRP	SpectraLink Radio Protocol
120	UTI	UTI
121	SMP	Simple Message Protocol
122	SM	SM
123	PTP	Performance Transparency Protocol
124	ISIS over IPv4	
125	FIRE	
126	CRTP	Combat Radio Transport Protocol
127	CRUDP	Combat Radio User Datagram
128	SSCOPMCE	
129	IPLT	
130	SPS	Secure Packet Shield
131	PIPE	Private IP Encapsulation within IP

#	Protocol	Protocol's Full Name
132	SCTP	Stream Control Transmission Protocol
133	FC	Fibre Channel
134	RSVP-E2E-IGNORE	
135	Mobility Header	
136	UDPLite	
137	MPLS-in-IP	
138	manet	
139	HIP	
140	Shim6	
141	WESP	
142	ROHC	
143 – 252	Unassigned	Unassigned
253		Use for experimentation and testing
254		Use for experimentation and testing
255	Reserved	

Further information can be found by researching RFC 5237.

Service groups

Just like some of the other firewall components, services can also be bundled into groups for ease of administration.

Creating a service group

1. Go to **Policy & Objects > Objects > Services**.
2. Select the down arrow next to **Create New**, select **Service Group**.
3. Input a Group Name for the list of services.
4. Input any additional information in the Comments field.
5. Choose a **Type**; either **Firewall** or **Explicit Proxy**.
6. Next to Members there is a dropdown menu that can be used to select from the available Services. It is possible to select more than 1 entry. Just select the green plus sign next to the field to add an additional entry. Select the "X"

icon in the field to remove an entry.

7. Press **OK**.

Example scenario: Supporting audio/visual conferencing



The feature, and the transmitting of data for the purpose of, Tele-conferencing or Audio/Visual Conferencing is covered by a number of standards:

- The IETF standard known as the Binary Floor Control Protocol (BFCP).
- RFC 4582, for SIP-based video devices
- The ITU standard H.239 (for H.323-based video devices)

While these standards have been set up by various authoritative bodies and can take place on different layers of the OSI model, they share common requirements that are addressed by the FortiGate firewall's ability to manage the traffic and the protocols involved. This means that the same ability that make the device RFC 4582 compliant makes it compliant with H.239 as well.

To demonstrate how services and service groups are used we show the setup of a firewall that will need to support the connectivity of a video conferencing unit. The FortiGate does not manipulate or change the content of the traffic but it does allow for the traffic to pass through the device. In this case it allow for only the needed traffic to pass through the device so as to allow the functionality of Audio Visual Conference call but not to allow other traffic through.

The theoretical location for this scenario is a hospital that hosts conferences and lectures from doctors from all over the world, sometimes from multiple locations, using video conferencing technology such as a Polycom Video Conference system. There is a special room set up with dedicated Ethernet connectivity to the Internet. A hospital has a lot of sensitive information going over its network so the setup has to be secure to prevent any chance of penetration.

The approach is fairly simple. The conference room has a dedicated port on the FortiGate (port #7) and its own LAN. We will assume that the interface has already been configured properly. Video conference traffic can come from the Internet to the Polycom in that room and traffic can get out to the Internet, but traffic going to other areas of the hospital network have to go through the FortiGate and traffic going from the Video Conference LAN is thoroughly filtered.

To give an idea of how extensive this can be, we will use an extreme case and include just about all of the services that could be commonly used in one of these setups. The protocols listed here may differ from other setups. It will depend on which features are being used and which equipment is within the network. Always check the documentation that comes with the set up before opening ports into your network.

VIP

In this particular case there is an IP address set aside for the conferencing system so a separate VIP is not needed for every port. One Virtual IP will be created for the system and then only the approved of protocols will be allows through the firewall.

Name	Vid-Conf_Room216
External Interface	wan1
External IP Address/Range	256.87.212.51 – 256.87.212.51
Mapped IP Address/Range	192.168.7.25 – 192.168.7.25
Port Forwarding	not selected

Creating an address for the subnet

In the same way that the VIP was created to identify and direct incoming traffic an address should be created to identify the addresses of computer that will be in the Conference room. This included computers on the LAN as well as the Teleconferencing equipment.

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
3. Fill out the fields with the following information:

Category	Address
Name	Port7_subnet
Type	Subnet
Subnet/IP Range	192.168.7.0/255.255.255.0
Interface	port7
Show in address list	checked

Configuring the services

Services already created:

The following are standard services that have already been created by default:

HTTP	TCP 80
SNMP	TCP 161-162/UDP 161-162
LDAP	TCP 389
HTTPS	TCP 443
SYSLOG	UDP 514

Existing services to be edited:

There are a few services that have already been created for you, but they need to be expanded to accommodate the list of protocols listed for this scenario.

The default h323 contains:

- TCP 1503
- UDP 1719
- TCP 1720

We need to add:

- TCP1719

The default SIP contains:

- UDP 5060

We need to add:

- TCP 5060

H323 service

1. Go to **Policy & Objects > Objects > Services**.
2. Scroll down to the section: **VoIP, Messaging & Other Applications**.
3. Select **H323**.
4. Select **Edit**.
5. In the Protocol section add the additional protocol:

Protocol Type	TCP
Destination port /Low	1719

6. Select **OK** to save.

SIP service

1. Go to **Policy & Objects > Objects > Services**.
2. Scroll down to the section: **VoIP, Messaging & Other Applications**.
3. Select **SIP**.
4. Select **Edit**.
5. In the Protocol section add the additional protocol:

Protocol Type	TCP
Destination port /Low	5060

6. Select **OK** to save.

Custom services that need to be created

There are a number of possible services that may need to be added from scratch rather than editing existing ones. While it is possible to create a single custom service that contains all of the open ports needed, it makes more sense to make this modular in case only a small subset of the service needs to be added to another policy.

Polycom API

1. Go to **Policy & Objects > Objects > Services**.
2. Select **Create New**.
3. Fill in the fields of the new service with the following information:

Name	Polycom API
Service Type	Firewall
Category	VoIP, Messaging & Other
Protocol Type	TCP/UDP/SCTP
Protocol	TCP/UDP/SCTP
Protocol	TCP
Destination Port - Low:	24
Destination Port - High:	<leave blank>

4. Select **OK**.

Polycom endpoints

1. Go to **Policy & Objects > Objects > Services**.
2. Select **Create New**.
3. Fill in the fields of the new service with the following information:

Name	Polycom Endpoints
Service Type	Firewall
Category	VoIP, Messaging & Other
Protocol Type	TCP/UDP/SCTP
Protocol	TCP
Destination - Low:	3230
Destination - High:	3253

4. Select **OK**.

Other Services to add in the same way:

Name of Service	Category	Protocol & Port #
LDAP secure communications	Authentication	TCP 636
Win 2000 ILS Registration	Network Services	TCP 1002
Gatekeeper discovery	VoIP, Messaging & Other Applications	TCP 1718
Audio Call Control	VoIP, Messaging & Other Applications	TCP 1731
Polycom proprietary Global directory data	VoIP, Messaging & Other Applications	TCP 3601
Polycom People+Content	VoIP, Messaging & Other Applications	TCP 5001
HTTP Server Push	Web Access	

Creating the service group

1. Go to **Firewall Objects > Service > Groups**.
2. Select **Create New**.
3. Build the Service group by filling in the fields with the following information

Group Name	A-V_Conference
Type	Firewall
Members (click in the drop down menu to add the following services)	<ul style="list-style-type: none"> • HTTP • SNMP • LDAP • HTTPS • SYSLOG • Polycom API • Polycom Endpoints • LDAP secure communications • Win 2000 ILS Registration • Gatekeeper discovery • Audio Call Control • Polycom proprietary Global directory data • Polycom People+Content • HTTP Server Push

Creating the IPS security profile

This is by no means the only way to set up this IPS filter, but it is the way that the fictional System Administrator wants it set up. Yours may be different.

1. Go to Security **Profiles > Intrusion Protection > IPS Sensors**.
2. Create a new sensor.

Name	A-V_Conference-incoming
-------------	-------------------------

3. Select **OK**.
4. In the newly created sensor, create a new IPS filter.

Sensor Type	Filter Based
Filter Options	Advanced
Severity	<ul style="list-style-type: none"> • Critical • High • Medium • Low
Target	Server
OS	Windows
Application	<ul style="list-style-type: none"> • IIS • other
Protocol Use the [Show more...] option	<ul style="list-style-type: none"> • HTTP • LDAP • SIP • SSL • H323
Packet logging	enabled

Based on these filters there should be somewhere in the neighborhood of 750 signatures that the FortiGate will run traffic against in the IPS engine.

Policies

Incoming policy

A policy has to be made to allow the traffic to come in from the Internet to connect to the Tele-conferencing server equipment.

1. Go to **Policy & Objects > Policy > IPv4**.
2. Select **Create New**.
3. Fill out the fields with the following information:

Policy Type	Firewall
Policy Subtype	Address
Incoming Interface	wan1
Source Address	all

Outgoing Interface	port7
Destination Address	Vid-Conf_Room216
Schedule	always
Service	A-V_Conference
Action	ACCEPT
Enable NAT	<not enabled>
Logging Options	Logging is a good idea but how much will depend on storage capabilities.
Security Profiles	Turn on IPS and choose "A-V_Conference-incoming"
Traffic Shaping, Web cache, WAN Optimization, Disclaimer:	The use of these features will depend on your network environment and should be decided by the network architect, as the decision will largely be based on network bandwidth, usage and importance of Video conferencing compared to other traffic.

4. Select **OK**.

The policy will then need to be put in the correct position in the sequence of the policies. Because it is a rather focused policy it should be acceptable to place it near the top of the policy order sequence.

Outgoing policy

A policy has to be made to allow the traffic to leave from the subnet in the conference room to the Internet, not only for the traffic for the Tele-conferencing equipment but for normal traffic of users on the Internet such as web research and email. The traffic is outgoing so there is less of a need for an Intrusion Protection System filter, but check with the network architect in case there is a need for using one of the other security profiles.

1. Go to **Policy & Objects > Policy > IPv4**.
2. Select **Create New**.
3. Fill out the fields with the following information:

Policy Type	Firewall
Policy Subtype	Address
Incoming Interface	port7
Source Address	Port7_subnet
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	any
Action	ACCEPT

Enable NAT	enabled Use Destination Interface Address
Logging Options	Logging is a good idea but how much will depend on storage capabilities.
Security Profiles	<see above>
Traffic Shaping, Web cache, WAN Optimization, Disclaimer:	The use of these features will depend on your network environment and should be decided by the network architect, as the decision will largely be based on network bandwidth, usage and importance of Video conferencing compared to other traffic.

4. Select OK.

The policy will then need to be put in the correct position in the sequence of the policies.

Firewall schedules

Firewall schedules control when policies are in effect. When you add a security policy on a FortiGate unit you need to set a schedule to determine the time frame in which that the policy will be functioning. While it is not set by default, the normal schedule would be always. This would mean that the policy that has been created is always function and always policing the traffic going through the FortiGate. The time component of the schedule is based on a 24 hour clock notation or military time as some people would say.

There are two types of schedules: One-time schedules and recurring schedules.

One-Time schedules are in effect only once for the period of time specified in the schedule. This can be useful for testing to limit how long a policy will be in effect in case it is not removed, or it can be used for isolated events such as a conference where you will only need a temporary infrastructure change for a few days.

The time frame for a One-time schedule is configured by using a start time which includes, Year | Month | Day | Hour | Minute and a Stop time which includes the same variables. So while the frequency of the schedule is only once it can last anywhere from 1 minute to multiple years.

Recurring schedules are in effect repeatedly at specified times of specified days of the week. The Recurring schedule is based on a repeating cycle of the days of the week as opposed to every x days or days of the month. This means that you can configure the schedule to be in effect on Tuesday, Thursday, and Saturday but not every 2 days or on odd numbered days of the month.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next.

Example

You want to schedule the use of Skype to only between noon (12:00) and 1 p.m. (13:00).

You could create a schedule that allows Skype traffic:

- Starting at Hour:12 and Minute: 00
- Stopping at Hour:13 and Minute: 00
- Set for days of the week: Sunday | Monday |Tuesday |Wednesday | Thursday | Friday | Saturday

Or you could have a schedule that blocks Skype traffic:

- Starting at Hour:13 and Minute: 00 (and goes to the next day)
- Stopping at Hour:12 and Minute: 00
- Set for days of the week: Sunday | Monday |Tuesday |Wednesday | Thursday | Friday | Saturday

Either way is effective for the task but other factors may make one method work better than another in certain situations of it could be just a preference in approach.

Schedule expiration

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as

the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the config firewall command enter the command:

```
set schedule-timeout enable
```

By default, this option is set to disable.

A few further settings are needed to make this work.

```
config firewall policy
  edit ID
    set firewall-session-dirty check-new
  end

config system settings
  set firewall-session-dirty check-policy-option
end
```

Firewall-session-dirty setting

The firewall-session-dirty setting has three options

check-all	CPU flushes all current sessions and re-evaluates them. [default]
check-new	CPU keeps existing sessions and applies policy changes to new sessions only. This reduces CPU load and the possibility of packet loss.
check-policy-option	Use the option selected in the firewall-session-dirty field of the firewall policy (check-all or check-new, as above, but per policy).

Schedule groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. The schedule parameter in the policy configuration does not allow for the entering of multiple schedules into a single policy so if you have a combination of time frames that you want to schedule the policy for then the best approach, rather than making multiple policies is to use a schedule group.

Example

Your Internet policy allows employees to visit Social Media sites from company computers but not during what is considered working hours. The offices are open a few hours before working hours and the doors are not locked until a few hours after official closing so work hours are from 9 to 5 with a lunch break from Noon to 1:00 p.m.

Your approach is to block the traffic between 9 and noon and between 1:00 p.m. and 5:00 p.m. This means you will need two schedules for a single policy and the schedule group handles this for you. Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

Schedule expiration

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the config firewall command enter the command:

```
set schedule-timeout enable
```

By default, this is set to `disable`.

Security policies

One of the foundations upon which a firewall works is the use of policies. These are what bring the other firewall objects and components together into an elegant mechanism for the governing of the traffic going through the network.

This Chapter includes information on the following topics:

- Firewall policies
- Security profiles
- SSL/SSH Inspection
- Identity Based Policies
- Device Identity Policies
- VPN Policies
- Interface Policies
- One-Arm IDS
- Local-In Policies
- Security Policy 0
- Deny Policies
- Accept Policies
- IPv6 Policies
- Fixed Port
- Endpoint Security
- Traffic Logging
- Quality of Service
- Policy Monitor

Firewall policies

The firewall policy is the axis around which most of the other features of the FortiGate firewall revolve. A large portion of the settings in the firewall at some point will end up relating to or being associated with the firewall policies and the traffic that they govern. Any traffic going through a FortiGate unit has to be associated with a policy. These policies are essentially discrete compartmentalized sets of instructions that control the traffic flow going through the firewall. These instructions control where the traffic goes, how it's processed, if it's processed and even whether or not it's allowed to pass through the FortiGate.

When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number). It also registers the incoming interface, the outgoing interface it will need to use and the time of day. Using this information the FortiGate firewall attempts to locate a security policy that matches the packet. If it finds a policy that matches the parameters it then looks at the action for that policy. If it is ACCEPT

the traffic is allowed to proceed to the next step. If the Action is DENY or a match cannot be found the traffic is not allowed to proceed.

The 2 basic actions at the initial connection are either ACCEPT or DENY:

- If the Action is ACCEPT, the policy action permits communication sessions. There may be other packet processing instructions, such as requiring authentication to use the policy. While you may not see it in the configuration there is the implied subset of the ACCEPT Action that include VPN policies, whether they be an IPsec VPN or SSL.
- If the Action is DENY, the policy action blocks communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped. A DENY security policy is needed when it is required to log the denied traffic, also called “violation traffic”.

The policy may contain a number of instructions for the FortiGate firewall in addition to the ACCEPT or DENY actions, some of which are optional. Instructions on how to process the traffic can also include such things as:

- Logging Traffic
- Authentication
- Network Address Translation or Port Address Translation
- Use Virtual IPs or IP Pools
- Caching
- Whether to use address or Identity based rules
- Whether to treat as regular traffic or VPN traffic
- What certificates to use
- Security profiles to apply
- Proxy Options
- Traffic Shaping

Firewall policy parameters

As mentioned before, for traffic to flow through the FortiGate firewall there must be a policy that matches its parameters:

Incoming interface

This is the interface or interfaces that the traffic is first connection to the FortiGate unit by. The exception being traffic that the FortiGate generates itself. This is not limited to the physical Ethernet ports found on the device. The incoming interface can also be a logical or virtual interface such as a VPN tunnel, a Virtual WAN link or a wireless interface.

Outgoing interface

After the firewall has processed the traffic it needs to leave a port to get to its destination and this will be the interface or interfaces that the traffic leaves by. This interface, like the **Incoming Interface** is not limited to only physical interfaces.

Source address

The addresses that a policy can receive traffic from can be wide open or tightly controlled. For a public webserver that the world at large should be able to access, the best choice will be “all”. If the destination is a private webserver that only the branch offices of a company should be able to access or a list of internal computers that

are the only ones allowed to access an external resource then a group of preconfigured addresses is the better strategy.

Additional parameters under the Source Address, though they are not mandatory are:

- **Source User(s)**

This parameter is based on a user identity that can be from a number of authentication authorities. It will be an account or group that has been set up in advance that can be selected from the dropdown menu. The exception to this is the feature that allows the importing of LDAP Users. When the feature is used, a small wizard window will appear to guide the user through the setup. The caveat is that the LDAP server object in the **User and Device > Authentication > LDAP Servers** section has to be already configured to allow the use of this import feature.

- **Source Device Type**

This parameter is for narrowing down the traffic sending devices to those that the FortiGate is familiar with. Again the contents of this parameter need to be a preconfigured object and these are defined at **User and Device > Device > Device Definitions**. This parameter can limit the devices that can connect to this policy to those specific MAC addresses that are already known by the FortiGate and are approved for the policy.

Destination address

In the same way that the source address may need to be limited, the destination address can be used as a traffic filter. When the traffic is destined for internal resources the specific address of the resource can be defined to better protect the other resources on the network. One of the specialized destination address options is to use a Virtual IP address. The destination address doesn't need to be internal you can define policies that are only for connecting to specific addresses on the Internet.

Schedule

The time frame that is applied to the policy. This can be something as simple as a time range that the sessions are allowed to start such as between 8:00 am and 5:00 pm. Something more complex like business hours that include a break for lunch and time of the session's initiation may need a schedule group because it will require multiple time ranges to make up the schedule.

Service

The service or service chosen here represent the TCP/IP suite port numbers that will most commonly be used to transport the named protocols or group of protocols. This will be a little different than Application Control which looks more closely at the packets to determine the actual protocol used to create them.

Without all six (possibly 8) of these things matching, the traffic will be declined. Each traffic flow requires a policy and the direction is important as well. Just because packets can go from point A to point B on port X does not mean that the traffic can flow from point B to point A on port X. A policy must be configured for each direction.

When designing a policy there is often reference to the traffic flow, but most communication is a two way connection so trying to determine the direction of the flow can be somewhat confusing. If traffic is HTTP web traffic the user sends a request to the web site, but most of the traffic flow will be coming from the web site to the user. Is the traffic flow considered to be from the user to the web site, the web site to the user or in both directions? For the purposes of determining the direction for a policy the important factor is the direction of the initiating communication. The user is sending a request to the web site so this is the initial communication and the web site is just responding to it so the traffic will be from the users network to the Internet.

A case where either side can initiate the communication like between two internal interfaces on the FortiGate unit would be a more likely situation to require a policy for each direction.

What is not expressly allowed is denied

One of the fundamental ideas that can be found in just about any firewall is the rule that anything that is not expressly allowed is by default denied. This is the foundation for any strategy of protecting your network. Right out of the box, once you have your FortiGate device connected into your network and hooked up with your ISP your network is protected. Nothing is getting out or in so it is not very convenient, but you don't have to worry that between the time you hooked it up and the point that you got all of the policies in place that someone could have gotten in and done something to your resources. The reason that this needs to be kept in mind when designing policies is because you cannot assume that any traffic will be allowed just because it makes sense to do so. If you want any kind of traffic to make it past the FortiGate firewall you need to create a policy that will allow that traffic. To maintain the protection of the network should also make sure that the any policy you create allows only the traffic you intend to go only to where you specifically want it to go and when you want it to go there.

Example

You have a web server on your network that is meant to provide a collaborative work environment web site for your employees and a partner company for a project over the course of the next 3 months.

It is theoretically possible to allow connections into your network to any device on that network for any service and at any time. The problem with this is that we might not want just anybody looking at those resources. Sadly, no matter how much it is wished otherwise, not everybody on the Internet can be trusted. Which means we now have to be very specific in our instructions as to what traffic to allow into the network. Each step that we take towards being more specific as to what we allow means that there is that much more that is not allowed and the level of protection of a resources is directly proportional to the amount of traffic that is not allowed. If somebody can't get at it they can't damage or steal it.

Limiting where the traffic is allowed to go to means that other computers on your network besides the web-server are protected.

- Limiting where the traffic is allowed to come from means that, if feasible, you can limit the systems that can access the web server to just employees or the partner company computers.
- Limiting the services to just web traffic means that a malicious person, even if they were connection from a computer at the partner organization could only use the features of web traffic to do anything malicious.
- Limiting the policy to the time span of the project would mean that even if the IT department forgot to remove the policy after the end of the project than no computer from the other company could be used to do anything malicious through the policy that allowed the traffic.

This is just a very basic example but it shows the underlying principles of how the idea that anything not expressly allowed is by default denied can be used to effectively protect your network.

Policy order

Another important factor in how firewall policies work is the concept of precedence of order or if you prefer a more recognizable term, "first come, first served".

It is highly likely that even after only a relatively small number of policies have been created that there will be some that overlap or are subsets of the parameters that the policies use to determine which policy should be matched against the incoming traffic. When this happens there has to be a method to determine which policy should be applied to the packet. The method which is used by most firewalls is based on the order of the sequence of the policies.

If all of the policies were placed in a sequential list the process to match up the packet would start at the top of the list and work its way down. It would compare information about the packet, specifically these points of information:

1. The interface the packet connected to the FortiGate firewall
2. The source address of the packet, and possibly the sending user and device.
3. The destination address of the packet
4. The interface the packet would need to use to get to the destination address based on the routing table
5. The port the packet is destined for
6. The time that the packet connected to the FortiGate

As soon as the a policy is reached that matches all of the applicable parameters, the instructions of that policy are applied and the search for any other matching policies is stopped. All subsequent policies are disregarded. Only 1 policy is applied to the packet.

If there is no matching policy among the policies that have been configured for traffic the packet finally drops down to what is always the last policy. It is an implicit policy. One of a few that are referred to by the term “policy0”. This policy denies everything.

The implicit policy is made up of the following settings:

- Incoming Interface: any
- Source Address: any
- Outgoing Interface: any
- Destination Address: any
- Action: DENY

The only setting that is editable in the implicit policy is the logging of violation traffic.

A logical best practice that comes from the knowledge of how this process works is to make sure that the more specific or specialized a policy is, the closer to the beginning of the sequence it should be. The more general a policy is the higher the likelihood that it could include in its range of parameters a more specifically targeted policy. The more specific a policy is, the higher the probability that there is a requirement for treating that traffic in a specific way.

Example

For security reasons there is no FTP traffic allowed out of a specific subnet so there is a policy that states that any traffic coming from that subnet is denied if the service is FTP, so the following policy was created:

Policy #1

Source Interface	Internal1
Source Address	192.168.1.0/24
Source User(s)	<left at default setting>
Source Device Type	<left at default setting>

Outgoing Interface	WAN1
Destination Address	0.0.0.0/0.0.0.0
Service	FTP
Schedule	always
Action	deny

Now as these things usually go it turns out that there has to be an exception to the rule. There is one very secure computer on the subnet that is allowed to use FTP and once the content has been checked it can then be distributed to the other computer on the subnet. So a second firewall policy is created.

Policy #2

Source Interface	Internal1
Source Address	192.168.1.38/32
Source User(s)	<left at default setting>
Source Device Type	<left at default setting>
Outgoing Interface	WAN1
Destination Address	0.0.0.0/0.0.0.0
Service	FTP
Schedule	always
Action	Allow

By default, a policy that has just been created will be placed last in the sequence so that it is less likely to interfere with existing policies before it can be moved to its intended position. If you look at Policy #2 you will notice that it is essentially the same as Policy #1 except for the Source Address and the Action. You will also notice that the Source Address of the Policy #2 is a subset of the Source address in policy #1. This means that if nothing further is done, Policy #2 will never see any traffic because the traffic will always be matched by Policy #1 and processed before it has a chance to reach the second policy in the sequence. For both policies to work as intended Policy #2 needs to be moved to before Policy #1 in the sequence.

Exception to policy order

There is a relevant exception to the normal policy order. Policies with VIPs don't appear to behave the same way. Traffic that is handled by VIPs is processed through the associated policy before the traffic is checked against other policies in the usual top down order.

This only appears inconsistent with the policy order rule because when handling the traffic, firewall policies are not the first thing checked. VIP translations are checked first, and if there is more than one VIP that the traffic fits, it is handled in the top down order that is followed by policies. If the traffic is not claimed by a policy in the VIP translation phase, it is checked against the routing rules. If it passes the routing checks, the traffic is allowed to be controlled by the policies.

This processing of traffic targeting a VIPs only applies if there is a policy that included the VIP and the traffic matches all of the criteria checks. There is no need to worry about creating VIPs that are not controlled by a policy.

There are security implications associated with this behavior. Administrators could assume that a policy will process traffic before it drops down to a policy with a VIP in it. This can allow traffic to pass through the firewall into a part of the network that it was not intended for, if it was to be allowed in at all. The way to prevent traffic being incorrectly allowed through a policy containing a VIP is to have that policy be more restrictive or to have a separate policy containing the same VIP deny the traffic earlier in the sequence.

As proof of the behaviour, look at the following traffic analysis of a packet sent to a VIP. You will see that the packet is translated even before it is allowed to pass through the firewall by a policy.

```
2015-06-10 06:33:21 id=20085 trace_id=1 func=print_pkt_detail line=4373 msg="vd-root
received a packet(proto=6, 24.114.222.34:51434->24.212.230.77:3389) from wan1. flag
[S], seq 1579917634, ack 0, win 8192"
2015-06-10 06:33:21 id=20085 trace_id=1 func=init_ip_session_common line=4522
msg="allocate a new session-01480894"
2015-06-10 06:33:21 id=20085 trace_id=1 func=fw_pre_route_handler line=174 msg="VIP-
10.10.66.2:3389, outdev-wan1"
2015-06-10 06:33:21 id=20085 trace_id=1 func=__ip_session_run_tuple line=2534 msg="DNAT
24.212.230.77:3389->10.10.66.2:3389"
2015-06-10 06:33:21 id=20085 trace_id=1 func=vf_ip4_route_input line=1596 msg="find a
route: flags=00000000 gw-10.10.66.2 via internal5"
2015-06-10 06:33:21 id=20085 trace_id=1 func=fw_forward_handler line=670 msg="Allowed by
Policy-15:"
```

Policy identification

When looking at the policy listing it can appear as if the policies are identified by the sequence number in the far left column. The problem is that this number changes as the position of the policy in the sequence changes. The column that correctly identifies the policy, and the value sticks with the policy is the "ID" column. This column is not shown by default in the listing but can be added to the displayed columns by right clicking on the column heading bar and selecting it from the list of possible columns.

When looking in the configuration file the sequence is based upon the order of the policies as they are in the file just as they are in the list in the GUI. However, if you need to edit the policy in the CLI you must use the ID number.

UUID support

Universally Unique Identifier (UUID) attributes have been added to policies to improve functionality when working with FortiManager or FortiAnalyser units. If required, the UUID can be set manually through the CLI.

CLI syntax:

```
config firewall {policy/policy6/policy46/policy64}
  edit 1
    set uuid <example uuid: 8289ef80-f879-51e2-20dd-fa62c5c51f44>
    next
  end
```

Security profiles

Where security policies provide the instructions to the FortiGate unit for controlling what traffic is allowed through the device, the Security profiles provide the screening that filters the content coming and going on the network. Security profiles enable you to instruct the FortiGate unit about what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A security profile is a group of options and filters that you can apply to one or more firewall policies. Security profiles can be used by more than one security policy. You can configure sets of security profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same security profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Security profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure security profiles in the Security Profiles menu and applied when creating a security policy by selecting the security profile type.

There is a separate handbook for the topic of the Security Profiles, but because the Security Profiles are applied through the Firewall policies it makes sense to have at least a basic idea of what the security profile do and how they integrate into the FortiGate's firewall policies. The following is a listing and a brief description of what the security profiles offer by way of functionality and how they can be configured into the firewall policies.

AntiVirus

Antivirus is used as a catch all term to describe the technology for protection against the transmission of malicious computer code sometimes referred to as malware. As anyone who has listened to the media has heard that the Internet can be a dangerous place filled with malware of various flavours. Currently, the malware that is most common in the Internet, in descending order, is Trojan horses, viruses, worms, adware, back door exploits, spyware and other variations. In recent years, not only has the volume of malicious software become greater than would have been believed when it first appeared but the level of sophistication has risen as well.

The Antivirus Filter works by inspecting the traffic that is about to be transmitted through the FortiGate. To increase the efficiency of effort it only inspects the traffic being transmitted via the protocols that it has been configured to check. Before the data moves across the FortiGate firewall from one interface to another it is checked for attributes or signatures that have been known to be associated with malware. If malware is detected, it is removed.

Web filtering

Malicious code is not the only thing to be wary of on the Internet. There is also the actual content. While the content will not damage or steal information from your computer there is still a number of reasons that would require protection from it.

In a setting where there are children or other sensitive people using the access provided by a connected computer there is a need to make sure that images or information that is not appropriate is not inadvertently displayed to them. Even if there is supervision, in the time it takes to recognize something that is inappropriate and then properly react can expose those we wish to protect. It is more efficient to make sure that the content cannot reach the screen in the first place.

In an organizational setting, there is still the expectation that organization will do what it can to prevent inappropriate content from getting onto the computer screens and thus provoking an Human Resources incident. There is also the potential loss of productivity that can take place if people have unfiltered access to the Internet. Some organizations prefer to limit the amount of distractions available to tempt their workers away from their duties.

The Web filter works primarily by looking at the destination location request for a HTTP(S) request made by the sending computer. If the URL is on a list that you have configured to list unwanted sites, the connection will be disallowed. If the site is part of a category of sites that you have configured to deny connections to the session will also be denied. You can also configure the content filter to check for specific key strings of data on the actual web site and if any of those strings of data appear the connection will not be allowed.

Application control

Application control is designed to allow you to determine what applications are operating on your network and to also filter the use of these applications as required. Application control is also for outgoing traffic to prevent the use of applications that are against an organization's policy from crossing the network gateway to other networks. An example of this would be the use of proxy servers to circumvent the restrictions put in place using the Web Filtering.

Intrusion protection (IPS)

Intrusion Prevention System is almost self explanatory. In the same way that there is malware out on the Internet that the network needs to be protected from there are also people out there that take a more targeted approach to malicious cyber activity. No operating system is perfect and new vulnerabilities are being discovered all of the time. An intrusion prevention system is designed to look for activity or behavior that is consistent with attacks against your network. When attack like behavior is detected it can either be dropped or just monitored depending on the approach that you would like to take.

As new vulnerabilities are discovered they can be added to the IPS database so that the protection is current.

Email filtering

Spam or unsolicited bulk email is said to account for approximately 90% of the email traffic on the Internet. Sorting through it is both time consuming and frustrating. By putting an email filter on policies that handle email traffic, the amount of spam that users have to deal with can be greatly reduced.

Data Leak Prevention (DLP)

Data Leak Prevention is used to prevent sensitive information from leaving your network. When people think of security in the cyber-world one of the most common images is that of a hacker penetrating your network and making off with your sensitive information, but the other way that you can lose sensitive data is if someone already on the inside of your network sends it out. This does not have to be an act of industrial espionage. It can just be a case of not knowing the policies of the organization or a lack of knowledge of security or laws concerning privacy.

For instance, a company may have a policy that they will not reveal anyone's Social Security number, but an employee emails a number of documents to another company that included a lengthy document that has a Social Security number buried deep within it. There is not malicious intent but if the information got out there could be repercussions.

If an organization has any information in a digital format that it cannot afford for financial or legal reasons, to leave its network, it makes sense to have Data Leak Prevention in place as an additional layer of protection.

VoIP

Voice over IP is essentially the protocols for transmitting voice or other multimedia communications over Internet Protocol networks such as the Internet. The Security Profiles VoIP options apply the SIP Application Level Gateway (ALG) to support SIP through the FortiGate unit. The SIP ALG can also be used to protect networks from SIP-based attacks.

ICAP

Internet Content Adaptation Protocol (ICAP) off loads HTTP traffic to another location for specialized processing. The purpose of this module when triggered is to send the incoming HTTP traffic over to a remote server to be processed thus taking some of the strain off of the resources of the FortiGate unit. The reasons for the specialized process could be anything from more sophisticated Antivirus to manipulation of the HTTP headers and URLs.

EndPoint control

EndPoint Control makes sure that certain standards are kept. When a computer on the Internet becomes connected to the FortiGate unit by VPN that computer is now part of the same network and therefore needs to be subject to the same levels of protection, not only to protect the computer but the network. In the EndPoint Control section you can set the minimum standards for things like AntiVirus software and VPN software.

Proxy option components

Any time a security profile that requires the use of a proxy is enabled the Proxy Options field will be displayed. Certain inspections defined in security profiles require that the traffic be held in proxy while the inspection is carried out and so the Proxy Options are there to define the parameters of how the traffic will be processed and to what level the traffic will be processed. In the same way that there can be multiple security profiles of a single type there can also be a number of unique Proxy Option profiles so that as the requirements for a policy differ from one policy to the next you can also configure a different Proxy Option profile for each individual policy or you can use one profile repeatedly.

The Proxy Options refer to the handling of the following protocols:

- HTTP
- SMTP
- POP3
- IMAP
- FTP
- NNTP
- MAPI
- DNS
- IM

The configuration for each of these protocols is handled separately.

It should also be noted that these configurations apply to only the Security Profiles Proxy-based processes and not the Flow-based processes.

The use of different proxy profiles and profile options

Just like other components of the FortiGate, there is the option for different Proxy Option profiles so that you can be very granular in your control of the workings of the FortiGate. In the case of the Proxy Option profiles the thing that you will want to focus on is the matching up of the correct profile to a firewall policy that is using the appropriate protocols. If you are creating a Proxy Option profile that is designed for policies that control SMTP traffic into your network you only want to configure the settings that apply to SMTP. You do not need or want to configure the HTTP components.

Oversized file log

This setting is for those that would like to log the occurrence of oversized files being processed. It does not change how they are processed it only enables the FortiGate unit to log that they were either blocked or allowed through. A common practice is to allow larger files through without antivirus processing. This allows you to get an idea of how often this happens and decide on whether or not to alter the settings relating to the treatment of oversized files.

The setting of the threshold for what is considered to be an oversized file is located in the Oversized File / Email Threshold that is found in some of the protocol options for the Proxy Options.

Protocol port mapping

While each of the protocols listed has a default TCP port that is commonly used, the level of granularity of control on the FortiGate firewall allows that the port used by the protocols can be individually modified in each separate Profile. It can also be set to inspect any port with flowing traffic for that particular protocol. The headers of the packets will indicate which protocol generated the packet. To optimize the resources of the unit the mapping and inspection of protocols can be enabled or disabled depending on your requirements.

Comfort Clients

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit begins scanning the file. During the buffering and scanning procedure, the user must wait. After the scan is completed, if no infection is found, the file is sent to the next step in the process flow. If the file is a large one this part of the process can take some time. In some cases enough time that some users may get impatient and cancel the download.

The comfort client feature mitigates this potential issue by feeding a trickle of data while waiting for the scan to complete so as to let the user know that processing is taking place and that there hasn't been a failure in the transmission. This slow transfer rate continues until the antivirus scan is complete. Once the file has been successfully scanned without any indication of viruses the transfer will proceed at full speed.

If there is evidence of an infection the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file. If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.



Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Oversized file/email threshold

This is another feature that is related to antivirus scanning. The FortiGate unit has a finite amount of resources that can be used to buffer and scan a file. If a large file such as an ISO image or video file was to be downloaded this could not only overwhelm the memory of the FortiGate, especially if there were other large files being downloaded at the same time, but could exceed it as well. For this reason, how to treat large files needs to be addressed.

A threshold is assigned to determine what should be considered an oversized file or email. This can be set at any size from 1 MB to 50 MB. Any file or email over this threshold will not be processed by the Antivirus Security Profiles. Once a file is determined to be oversized it must be then determined whether to allow it or to block it.

These settings are not a technical decision but a policy one that will depend on your comfort level with letting files into your network. As there often is, there is a compromise between convenience or ease of use and security. If you want to go for a high peace of mind level you can configure the firewall to block oversized files and thus no files would be coming into the network that have not been scanned. If you are looking for optimizing the memory of the FortiGate unit and making sure that everybody is getting the files they want, you can lower the threshold and allow files that are over the threshold.



It should be noted that in terms of probability that malware is more likely to be found in smaller files than in larger files. A number of administrators take this into account when they lower the default threshold so as to lessen the impact on memory if they see the FortiGate unit going into conserve mode on a regular basis.

Chunked bypass

The HTTP section allows the enabling of "Chunked Bypass". This refers to the mechanism in version 1.1 of HTTP that allows a web server to start sending chunks of dynamically generated output in response to a request before actually knowing the actual size of the content. Where dynamically generated content is concerned this means

that there is a faster initial response to HTTP requests. From a security stand point it means that the content will not be held in the proxy as an entire file before proceeding.

Allow fragmented messages

The specifications of RFC 2046 allow for the breaking up of emails and sending the fragments in parallel to be rebuilt and read at the other end by the mail server. It was originally designed to increase the performance over slower connections where larger email messages were involved. It will depend on your mail configuration if this is even possible for your network but outside of Microsoft Outlook and Outlook Express, not many email clients are set up to break up messages like this. The drawback of allowing this feature is that if malware is broken up between multiple fragments of the message the risk is run that it will not be detected by some antivirus configurations because the code may not all be present at the same time to identify.

Append email signature

The Append Email Signature is used when an organization would like to ensure that over and above our in this case underneath the existing personal signatures of the sender, all of the emails going out of their network have the appropriate “boilerplate”, for lack of a better term. These appended emails do not replace existing signatures. They are as the feature states, appended to the email.

Examples could include things like:

- Without prior approval the email should not be forwarded.
- Please be environmentally friendly and don't print out emails
- For questions regarding the purchasing of our products please call...

It can be anything that the organization would like as long as it is in text format. The use of this feature usually works best in an environment where there is some standardization of what goes into the personal signatures of the senders so that there is no duplication or contradiction of information in the signatures.

SSL/SSH inspection

While the profile configuration for this is not found in the Security Profiles section but in the Policy Section, it is set in the policy along with the security profiles. This sort of analysis is some times referred to as deep scanning.

Deep Inspection works along the following lines. If your FortiGate unit has the correct chipset it will be able to scan SSL encrypted traffic in the same way that regular traffic can be scanned. The FortiGate firewall will essentially receive the traffic on behalf of the client and open up the encrypted traffic. Once it is finished it re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack. By enabling this feature, it allows the FortiGate firewall to filter on traffic that is using the SSL encrypted protocol.

The encrypted protocols that can be inspected are:

- HTTPS
- SMTPS
- POP3S
- IMAPS
- FTPS

Before the invention of SSL inspection, scanning regular web traffic can be circumvented by using the prefix https:// instead of http:// in the URL. SSL inspection prevents this circumvention. However, because when the

encrypted traffic is decrypted it has to be re-encrypted with the FortiGate's certificate rather than the original certificate it can cause errors because the name on the certificate does not match the name on the web site.

At one point deep inspection was something that was either turned on or off. Now individual deep inspection profiles can be created depending on the requirements of the policy. Depending on the Inspection Profile, you can:

- Configure which CA certificate will be used to decrypt the SSL encrypted traffic.
- Configure which SSL protocols will be inspected.
- Configure which ports will be associated with which SSL protocols for the purpose of inspection.
- Configure which websites will be exempt from SSL inspection
- Configure whether or not to allow invalid SSL certificates.
- Configure whether or not SSH traffic will be inspected.

Inspection exemption

When you are using a browser to visit SSL encrypted sites and we are using a certificate that does not match the certificate of the site, we are presented with a warning message and the option of continuing, using the untrusted certificate, or terminating the session. However, there are a number of applications that use SSL encrypted traffic. If the application detects SSL traffic that wasn't signed with a certificate that it trusts it will not allow the traffic. The applications do not give the option to manually indicate that we trust the certificate or the site.

If the option is available, the customer may choose to import needed SSL certificates into Local Certificates and configure a policy for communication for that application.

The assist in preventing loss of access to these site but still enabling the SSL inspection of the rest of the internet traffic, a method of exempting either Website categories or specific sites has been developed. To exempt a large group of sites the profile can be configure to exempt FortiGuard Categories. There are 3 of these categories preselected due to the high likelihood of issues with associated applications with the type of websites included in these categories.

- Health and Wellness
- Personal Privacy
- Finance and Banking

Other more specific websites can be added to the exemption list by creating addresses for them at **Policy & Objects > Objects > Addresses**. The adding of addresses is done by selection from a drop down menu. There is an option at the bottom of the list to create a new address, but otherwise only preconfigured addresses that are configured to be on the "Any" interface will be available for selection.

Examples of sites that you may want to configure for exemption so that there will be no interference due to certificate issues:

Apple

- *.appstore.com
- *.apple.com
- *.itunes.apple.com
- *.icloud.com
- swscan.apple.com

Dropbox

- *.dropbox.com

Skype

- *.messenger.live.com

Windows Updates

- update.microsoft.com

Allow invalid SSL certificate

This setting was something that used to be part of the **Proxy Options**, but now that SSL inspection has its own configuration setting it is configured with those. It might seem like a straight forward decision that the allowing of invalid SSL certificates must be bad and therefore should not be allowed, but there can be some reasons that should be considered. The issues at hand are the reasons to use a SSL certificate and the reasons that a certificate will be considered invalid.

At a purely technical level, a properly formed certificate will encrypt the data so that it can only be read by the intended parties and not be read by anyone sniffing traffic on the network. For this reason, people will often use self-signed certificates. These self signed certificates are free and will encrypt the data just as well as those purchased from any of the big vendors of certificates, but if they are not listed as an approved Certificate Authority (CA) the certificates will be considered invalid.

On the other hand, one of the services the vendors provide is verification of identity of those that purchase their certificates. This means that if you see a valid certificate from a site that identified itself as being from "valid-company.com" that you can be reasonably sure that the site does belong to that company and not a false site masquerading as being part of that company.

Creating or editing an SSL/SSH inspection profile

1. Go to **Policy & Objects > Policy > SSL/SSH Inspection**.

This will open to one of the existing profiles.

The links for the actions are located in the upper right hand corner of the window.

- To view a list of the existing profiles select the List icon (a page) at the far right.
- To clone an existing profile, select the Clone icon (one page behind another), second from the right
- To create a new profile, select the Create New icon ("+" symbol), third from the right.
- To view or edit an existing profile, choose it from the dropdown menu field.

2. **Name Field:**

Give the Profile an easily identifiable name that references its intent.

3. **Comments Field:**

Enter any additional information that might be needed by administrators, as a reminder of the profile's purpose and scope.

4. **SSL Inspection Options:**

a. Enable SSL Inspection of:

- Multiple Clients Connecting to Multiple Servers - Use this option for generic policies where the destination is unknown.

- Protecting SSL Server - Use this option when setting up a profile customized for a specific SSL server with a specific certificate.
 - b. CA Certificate
Use the drop down menu to choose which one of the installed certificates to use for the inspection of the packets.
 - c. Inspection Method
The options here are:
 - SSL Certificate Inspection - only inspects the certificate, not the contents of the traffic.
 - Full SSL Inspection - inspects all of the traffic.
 - d. Inspect All Ports
Enable the ability to inspect all ports by checking the box. If the feature is not enabled, specify in the field next to the listed protocols, the port through which that protocols traffic will be inspected. Traffic of that protocol going through any other port will not be inspected.
5. **Exempt from SSL Inspection:**
Use the dropdown menus in this section to specify either a FortiGuard Web Category or addresses that will be exempt from SSL inspection.
- a. Web Categories
By default the categories of Health and Wellness, Personal Privacy, and Finance and Banking have been added as these are one that are most likely to have applications that will require a specific certificate.
 - b. Addresses
These can be any of the Address objects that have an interface of "Any".
6. **SSH Inspection Options:**
- a. SSH Deep Scan
Toggle the grey on button so that it is:
Greyed out to disable the feature
Opaque and vibrate to enable the feature
 - b. SSH Port
The available options are:
 - **Any** - choosing this option will search all of the traffic regardless of service or TCP/IP port for packets that conform to the SSH protocol
 - **Specify** - choosing this option will restrict the search for SSH protocol packets to the TCP/IP port number specified in the field. This is not as comprehensive but it is easier on the performance of the firewall.
 - d. Protocol Actions
 - Exec - Block, Log or neither. Select using check boxes.
 - Port-Forward - Block, Log or neither. Select using check boxes.
 - SSH-Shell - Block, Log or neither. Select using check boxes.
 - X11-Filter - Block, Log or neither. Select using check boxes.
7. **Common Options:**
- a. Allow Invalid SSL Certificates
Check the box to enable the passing of traffic with invalid certificate
 - b. Log Invalid Certificates
Check the box to have the Logging function record traffic sessions that contained invalid certificates



The **Enable SSH Deep Scan** feature is enabled by default when creating a new SSL/SSH Inspection profile. There are situations where this feature can cause issues so be sure that you would like it enabled before applying it.



The context location for configuring the SSL/SSH Inspection in the CLI is:

```
config firewall ssl-ssh-profile
```

Viewing firewall policies

When you first go into the Policy window, found by going to Policy > Policy > Policy, you will see a table with a menu bar across the top. The menu bar will have the following items:

At the top left:

- Create New (with a “+” sign on the left and a downward pointing triangle on the right)
- Clone
- Delete
- Column Settings
- Filter Settings

At the top right:

- Section View
- Global View

The items at the top right with their radio buttons represent the 2 potential views that the policies can be displayed in.

The Global View shows all of the policies in the order of their sequence. With the default settings you will be able to see the sequence number in a column close to the left side of the table.

The Section view is similar to the Global View except that as the name implies it is divided into sections. By default the sections are based on the paths between the interfaces. These can be referred to as “interface pairings”. For instance, all of the policies referencing traffic from WAN1 to DMZ will be in one section. The policies referencing traffic from DMZ to WAN1 will be in another section.

The sections are collapsible so that you only need to look at the sections with policies you are interested in. It is possible to add customized subsections within the default sections of interface pairings. This would be useful in a situation where you have a lot of policies and would like to further compartmentalize them by common attributes so that things are easier to find.

The default column headings are:

- [Check box icon]
- Seq.#
- Source
- Destination
- Authentication
- Schedule

- Service
- Action
- Log

The columns that are shown are configurable. All but the first 2 can be removed or their position changed. There are also a number of other columns that display information about the policies that can be added. One of the more useful ones that can be added is the ID column. The reason for adding this one is that policies are referenced by their ID number for simplicity and ease of administration. If you are looking in the CLI you will see that the only designation for a policy is its number and if you wish to change the order of a policy you will be asked to move it before or after another policy by referencing its number.

How “Any” policy can remove the section view

The FortiGate unit will automatically change the view on the policy list page to Global View whenever a policy containing “any” in the Source interface/zone or Destination interface/zone is created. If the Section View is greyed out it is likely that one or more of the policies has “any” as a Source or Destination interface.

With the use of the “any” the policy should go into multiple sections because it could effectively be any of a number of interface pairings. As mentioned, policies are sectioned by using the interface pairings (for example, port1 -> port2) and each section has its own specific policy order. The order in which a policy is checked for matching criteria to a packet’s information is based solely on the position of the policy within its section or within the entire list of policies as a whole but if the policy is in multiple sections at the same time there is no mechanism for placing the policy in a proper order within all of those sections at the same time because it is a manual process and there is no parameter to compare the precedence of one section or policy over the other. Thus a conflict is created. In order to resolve the conflict the FortiGate firewall removes that aspect of the sections so that there is no need to compare and find precedence between the sections and it therefore has only the Global View to work with.

Security policy configuration extensions

When first creating the policy the configuration form will ask for a choice between the policy types of Firewall or VPN, Firewall being the default. Choosing whether or not to leave the selection as Firewall is straight forward. If the policy is not a policy based VPN policy then it is a Firewall policy type.

There are essentially 2 types of VPN connections, Interface Based and Policy Based. In an Interface Based VPN tunnel a logical interface is created that can be seen as an interface by the policies in the same way that any of the physical interfaces can be seen. Therefore to govern the traffic a regular policy will work. The policy based VPN tunnels work slightly different and therefore need a slightly different policy configuration. For a more detail explanation of the difference between the types of VPN tunnels refer to the VPN documentation found in the VPN handbooks or in the VPN section of the Complete Administration Guide.

Once either the Firewall or the VPN type has been chosen there is then a choice between one of subtypes for each of the Policy types. For the Firewall type of policy the subtypes are:

- Address
- User Identity
- Device Identity

The Address subtype refers to policies where access through the FortiGate firewall is dependant on the source location of the addresses of the devices involved in the traffic matched to the policy.

The User Identity subtype refers to policies where access through the FortiGate firewall is dependant on the users credentials or Identity.

The Device Identity subtype refers to policy where access through the FortiGate firewall is dependant on the specific device being used based on the MAC address of the device or belonging to a group of devices that are based on device types or belonging to custom made groups.

For the VPN type the subtypes are:

- IPsec
- SSL-VPN

As expected the two subtypes are the two different types of VPN tunnels that the FortiGate firewall supports in a policy based configuration.

Identity-based policies

Identity based policies are ones in which there is the additional component of either an account identity or device identity. The inclusion of one or both of these components adds an extra dimension of complexity to working with these policies in the context of the other policies so while the extra security and granularity of control are beneficial, extra care must be taken when configuring the policies themselves and how they are positioned in the policy sequence. The actual configuration of these identities are explained in detail in the Authentication Handbook.

Identity-based security policies are usually configured for IPsec or SSL VPN traffic since this type of traffic usually requires authentication from network users.

Identity-based policy positioning

In non-identity based policies, if non of the 6 mandatory policy parameters matches the header of the traffic packets the parameters are compared against the next policy in sequence. Because those parameters are mandatory there is always a value to test against and whether or not the policy applies is certain. The fact that the identity parameters are not required makes knowing whether or not the correct policy will be applied less obvious.

Originally, the identity aspect of a policy was an entire sub-policy checking sequence within each policy, including its own 0 policy at the end of the sequence. If all of the other parameters match the policy would then compare the traffic's identity with the list of identity groups in the policy starting at the beginning of the sequence and going through them until an identity was found that matched and then the rules for that identity group would be applied. If the traffic's identity did not match any of those listed in the policy it go to the last identity in the policy would be everyone and the Action would be deny.

The identity aspects of policies have now been incorporated in a single flat configuration that makes them a fundamental part of the policy rather than something that is added to the policy. This is simpler and allows for more complex combinations of address identification, user authentication and device determination that were not possible with previous policy configurations. Both user groups and device groups can be part of the same policy. Because the identity aspects are optional, more flexibility in creating policies that use authentication is possible.

Identity fall through rules

The fall through rules for policies in 5.2 have changed so that they are more in keeping with the practices of other vendors. This makes it easier for users used to other firewalls to configure the policies and it also makes it simpler to convert the policies of other firewalls to be used on a FortiGate firewall.

Previously, if traffic reached an identity policy and the user or device was not a member of one of the groups specified it would fall through to the implicit deny all policy. This meant that any traffic that reached that policy would have to be authenticated and a member of one of the listed groups. If the 6 required parameters matched, the traffic would not be getting past this policy.

The approach is now to treat the the identity parameters, if they exist, the same as the other parameters, in that if they do not match any listed in the policy, the traffic drops down to the next policy.

Example:

There are three policies where all the parameters are the same except:

- Policy # 1 - Source User Group A is assigned profile A
- Policy # 2 - Source User Group B s assigned profile B
- Policy # 3 - Source User(s) and Source Device Type are empty

Traffic that matches all of the required parameters will be processed as follows:

- Traffic authenticated as being from User Group A will be processed by Policy # 1.
- Traffic authenticated as being from User Group B will be processed by Policy # 2.
- Traffic with no authenticated users will be processed by Policy # 3.
- Traffic authenticated as being from User Group C will be processed by Policy # 3.

In the methodology before FortiOS 5.2, traffic authenticated as being User Group B, User Group C or no authenticated user at all would have been stopped at Policy # 1.

The CLI command “fall-through-unauthenticated” that was added in 5.0.1 attempted to allow a process similar to this, but only applied to unauthenticated traffic and not authenticated traffic that didn’t match the list of groups is the sub-policy. The current methodology is not subject to the same limitation and alleviates the need for the function of this command so the command has been removed from the CLI.

Implicit protocols

In previous versions of the firmware, the protocols that were used to authenticate such as HTTP, HTTPS, FTP, and Telnet, were supported on the policy whether or not they were included in the supported services. In 5.2, the protocol needed to authenticate needs to be included in the list of allowed services in order the the authentication to take place.

For example, if you have a VIP coming into your network that is for connecting to some security webcams located in your data center that use custom services or ports to connect to, if you are using an identity policy you would also have to include HTTP or HTTPS in the services list in order to actually authenticate.



Another formerly implicit protocol that is not supported automatically in 5.2 is port 53 (DNS). If you are limiting the services of a protocol to web based protocols such as HTTP or HTTPS don't forget to add DNS so that the domain names can be resolved.



When upgrading the firmware from version 5.0.x to 5.2.x, a policy with either an identity or device sub-policy will automatically convert from a single policy with sub-policies to a separate policy for each identity based sub-policy.

VPN policies

At one point, if you wanted to have secure digital communications between 2 points a private network would be created. This network would only allow the people that were intended to get the communications on it. This is very straightforward if the 2 points are in the same room or even in the same building. It can all be done physically. If you are supposed to be on the secure network

VPNs are an answer to one of today's biggest concerns, how to make digital communications secure between to points that must communicate over the Internet which anybody can have access to

IPsec policies

IPsec policies allow IPsec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate firewall interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate firewall interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.

For a route-based (interface mode) VPN, you do not configure an IPsec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPsec virtual interface as the source or destination interface, as appropriate.

Interface policies

Interface policies are implemented before the "security" policies and are only flow based.

This feature allows you to attach a set of IPS policies with the interface instead of the forwarding path, so packets can be delivered to IPS before entering firewall. This feature is used for following IPS deployments:

- One-Arm: by defining interface policies with IPS and DoS anomaly checks and enabling sniff-mode on the interface, the interface can be used for one-arm IDS;
- IPv6 IPS: IPS inspection can be enabled through interface IPv6 policy. Only IPS signature scan is supported in FortiOS 5.2. IPv6 DoS protection is not supported;
- Scan traffics that destined to FortiGate;
- Scan and log traffics that are silently dropped or flooded by Firewall or Multicast traffic.

IPS sensors can be assigned to an interface policy. Both incoming and outgoing packets are inspected by IPS sensor (signature).

Here is an example of an interface policy,

```
config firewall interface-policy
edit 1
    set status enable
    set interface "port14"
    set srcaddr "all"
    set dstaddr "all"
```

```
set service "ALL"
set application-list-status disable
set ips-sensor-status enable
set ips-sensor "default"
set av-profile-status disable
set webfilter-profile-status disable
set spamfilter-profile-status disable
set dlp-sensor-status disable
set label "Port 14 Interface Policy"
next
end
```

DoS protection

Denial of Service (DoS) policies are primarily used to apply DoS anomaly checks to network traffic based on the FortiGate interface it is entering as well as the source and destination addresses. DoS checks are a traffic anomaly detection feature to identify network traffic that does not fit known or common traffic patterns and behavior. A common example of anomalous traffic is the denial of service attack. A denial of service occurs when an attacking system starts an abnormally large number of sessions with a target system. The large number of sessions slows down or disables the target system, so that legitimate users can no longer use it.

DoS policies are similar to firewall policies except that instead of defining the way traffic is allowed to flow, they keep track of certain traffic patterns and attributes and will stop traffic displaying those attributes. Further, DoS policies affect only incoming traffic on a single interface. You can further limit a DoS policy by source address, destination address, and service.

DoS configurations have been changed a couple of times in the past. In FortiOS 4.0, DoS protection was moved to the interface policy, so when it is enabled, it is the first thing checked when a packet enters FortiGate. Because of this early detection, DoS policies are a very efficient defence that uses few resources. Denial of service attacks, for example, are detected and its packets dropped before requiring security policy look-ups, antivirus scans, and other protective but resource-intensive operations.

A DoS policy examines network traffic arriving at an interface for anomalous patterns usually indicating an attack. This does not mean that all anomalies experience by the firewall are the result of an intentional attack.

Because an improperly configured DoS anomaly check can interfere with network traffic, no DoS checks are preconfigured on a factory default FortiGate unit. You must create your own before they will take effect. Thresholds for newly created sensors are preset with recommended values that you can adjust to meet the needs of your network.

To create a Denial of Service policy determine if it needs to be an IPv4 or IPv6 policy, then goto:

Policy & Objects > Policy > DoS Policy for IPv4.

Policy & Objects > Policy > IPv6 DoS Policy for IPv6.



It is important to know normal and expected network traffic before changing the default anomaly thresholds. Setting the thresholds too low could cause false positives, and setting the thresholds too high could allow otherwise avoidable attacks.

Settings used in configuring DoS

Incoming interface

The interface to which this security policy applies. It will be the that the traffic is coming into the firewall on.

Source address

This will be the address that the traffic is coming from and must be a address listed in the Address section of the Firewall Objects. This can include the predefined “all” address which covers any address coming in on any interface. Multiple addresses or address groups can be chosen

Destination address

This will be the address that the traffic is addressed to. In this case it must be an address that is associated with the firewall itself. For instance it could be one of the interface address of the firewall, a secondary IP address or the interface address assigned to a Virtual IP address. Just like with the Source Address this address must be already configured before being used in the DoS policy. Multiple addresses, virtual IPs or virtual IP groups can be chosen.

Service

While the Service field allows for the use of the ALL service some administrators prefer to optimize the resources of the firewall and only check on the services that will be answered on an interface. Multiple services or service groups can be chosen.

Anomalies

The anomalies can not be configured by the user. They are predefined sensors set up for specific patterns of anomalous traffic

The anomalies that have been predefined for use in the DoS Policies are:

Anomaly Name	Description	Recommended Threshold
tcp_syn_flood	If the SYN packet rate of new TCP connections, including retransmission, to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
tcp_port_scan	If the SYN packet rate of new TCP connections, including retransmission, from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second.
tcp_src_session	If the number of concurrent TCP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.

Anomaly Name	Description	Recommended Threshold
tcp_dst_session	If the number of concurrent TCP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
udp_flood	If the UDP traffic to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
udp_scan	If the number of UDP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
udp_src_session	If the number of concurrent UDP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
udp_dst_session	If the number of concurrent UDP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
icmp_flood	If the number of ICMP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	250 packets per second.
icmp_sweep	If the number of ICMP packets originating from one source IP address exceeds the configured threshold value, the action is executed.	100 packets per second.
icmp_src_session	If the number of concurrent ICMP connections from one source IP address exceeds the configured threshold value, the action is executed.	300 concurrent sessions
icmp_dst_session	If the number of concurrent ICMP connections to one destination IP address exceeds the configured threshold value, the action is executed.	3000 concurrent sessions
ip_src_session	If the number of concurrent IP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
ip_dst_session	If the number of concurrent IP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
sctp_flood	If the number of SCTP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second

Anomaly Name	Description	Recommended Threshold
sctp_scan	If the number of SCTP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second
sctp_src_session	If the number of concurrent SCTP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions
sctp_dst_session	If the number of concurrent SCTP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions

Status

The status field is enabled to enable the sensor for the associated anomaly. In terms of actions performed there is no difference between disabling a sensor and having the action as “Pass” but by disabling sensors that are not being used for blocking or logging you can save some resources of the firewall that can be better used elsewhere.

Logging

Regardless of whether the traffic is blocked or passed through the anomalous traffic will be logged.

Pass

Allows the anomalous traffic to pass through unimpeded.

Block

For Thresholds based on the number of concurrent sessions blocking the anomaly will not allow more than the number of concurrent sessions set as the threshold.

For rate based thresholds where the threshold is measured in packets per second, the Action setting “Block” prevents the overwhelming of the firewall by anomalous traffic in one of 2 ways. Setting which of those 2 ways will be issued is determined in the CLI.

- continuous - blocks any packets that match the anomaly criteria once the threshold has been reached
- periodical - allows matching anomalous traffic up to the rate set by the threshold.



If the period for a particular anomaly is 60 seconds, such as those where the threshold is measured in concurrent sessions, after the 60 second timer has expired, the number of allowed packets that match the anomaly criteria is reset to zero. This means that if you allow 10 sessions through before blocking, after the 60 seconds is up, another 10 will be allowed. The attrition of sessions from expiration should keep the allowed sessions from reaching the maximum.

To set the type of block action for the rate based anomaly sensors:

```
config ips global
    set anomaly-mode continuous
    set anomaly-mode periodical
end
```

Threshold

The threshold can be either in terms of concurrent session or in packets per second depending on which sensor is being referred to.

One-arm IDS

Interface-based policy only defines what and how IPS functions are applied to the packets transmitted by the interface. It works no matter if the port is used in a forwarding path or used as an One-Arm device.

To enable One-Arm IDS, the user should first enable sniff-mode on the interface,

```
config system interface
  edit port2
    set ips-sniffer-mode enable
  next
end
```

Once sniff-mode is turned on, both incoming and outgoing packets will be dropped after IPS inspections. The port can be connected to a hub or a switch's SPAN port. Any packet picked up by the interface will still follow the interface policy so different IPS and DoS anomaly checks can be applied.

IPv6 IPS

IPv6 IPS signature scan can be enabled by interface policy. The user can create an normal IPS sensor and assign it to the IPv6 interface policy.

```
config firewall interface-policy6
  edit 1
    set interface "port1"
    set srcaddr6 "all"
    set dstaddr6 "all"
    set service6 "ANY"
    set ips-sensor-status enable
    set ips-sensor "all_default"
  next
end
```

Traffic destined to the FortiGate unit

IPS enabled in firewall policies can only inspect the traffic pass through FortiGate unit, not the traffic destined to FortiGate unit. Enabling IPS in interface-policy allows IPS to pick up any packet on the interface so it is able to inspect attacks targeting FGT.

Dropped, flooded, broadcast, multicast and L2 packets

In many evaluation or certification tests, FortiGate firewall is often required to log any packets dropped by the firewall. In most of cases, these packets are of invalid headers so firewall just drops them silently. It is natural to forward all these packets to IPS first so FortiGate firewall is able to generate logs for invalid packets.

Flooded, broadcast and multicast traffics do not reach any of services in the forwarding path. They can be inspected by the interface policy as long as they match the addresses defined. Potentially, L2 packets can also be sent to IPS for inspection through interface-policy, but it is not enabled in FortiOS 5.2.

GUI and CLI

Now in FortiGate, there are two places that IPS can be enabled, in a firewall policy and in an interface policy. In the firewall policy implementation, IPS sensor can be configured in both CLI and GUI. When adding an IPS sensor to an interface policy it must be done through the CLI. There is no GUI input window for the “Interface Policy”. There is however, a DoS Policy section in the GUI.

Local-In policies

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
  edit <policy_number>
    set intf <source_interface>
    set srcaddr <source_address>
    set dstaddr <destination_address>
    set action {accept | deny}
    set service <service name>
    set schedule <schedule_name>
  end
```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12, represented by the address object mgmt-comp1, using SSH on port 3 (192.168.21.77 represented by the address object FG-port3) using the Weekend schedule which defines the time the of access.

```
config firewall local-in-policy
  edit <1>
    set intf port3
    set srcaddr mgmt-comp1
```

```

set dstaddr FG-port3
set action accept
set service SSH
set schedule Weekend
end

```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```

config firewall local-in-policy
edit <policy_number>
set status disable
end

```

Use the same commands with a status of enable to use the policy again.

Local-in policies are also supported for IPv6 by entering the command

```
config firewall local-in-policy6.
```

Security policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPsec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPsec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate firewall logs, you may find a log field entry indicating policyid=0. The following log message example indicates the log field policyid=0 in bold.

```

2008-10-06 00:13:49 log_id=0022013001 type=traffic subtype=violation pri=warning
vd=root SN=179089 duration=0 user=N/A group=N/A rule=0 policyid=0 proto=17
service=137/udp app_type=N/A status=deny src=10.181.77.73 srcname=10.181.77.73
dst=10.128.1.161 dstname=10.128.1.161 src_int=N/A dst_int="Internal" sent=0 rcvd=0
src_port=137 dst_port=137 vpn=N/A tran_ip=0.0.0.0 tran_port=0

```

Deny policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.



There is a disparity in the effectiveness of deny policies. Only deny policies that contain VIPs will block traffic directed at those VIPs. Policies with VIPs are processed before other policies, so using a deny policy earlier in the list will not work. For more on this topic, read the ["Exception to policy order" on page 93](#).

Accept policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPsec VPN.

Fixed port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable fixedport when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
  edit <policy-id>
    ...
    set fixedport enable
    ...
  end
```

However, enabling fixedport means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Endpoint security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance

and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed. Currently, FortiClient Endpoint Security is available for Microsoft Windows 2000 and later only.

For more information about endpoint security, see the Security Profiles chapter in the FortiOS Handbook.

Traffic logging

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Depending on what the FortiGate unit has in the way of resources, there may be advantages in optimizing the amount of logging taking places. This is why in each policy you are given 3 options for the logging:

- **No Log** - Does not record any log messages about traffic accepted by this policy.
- **Log Security Events** - records only log messages relating to security events caused by traffic accepted by this policy.
- **Log all Sessions** - records all log messages relating to all of the traffic accepted by this policy.

Depending on the the model, if the Log all Sessions option is selected there may be 2 additional options. These options are normally available in the GUI on the higher end models such as the FortiGate 600C or larger.

- **Generate Logs when Session Starts**
- **Capture Packets**

You can also use the CLI to enter the following command to write a log message when a session starts:

```
config firewall policy
  edit <policy-index>
    set logtraffic-start
  end
```

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13
05:23:47
log_id=4
type=traffic
subtype=other
pri=notice
vd=root
status="start"
src="10.41.101.20"
srcname="10.41.101.20"
src_port=58115
dst="172.20.120.100"
dstname="172.20.120.100"
dst_country="N/A"
dst_port=137
```

```

tran_ip="N/A"
tran_port=0
tran_sip="10.31.101.41"
tran_sport=58115
service="137/udp"
proto=17
app_type="N/A"
duration=0
rule=1
policyid=1
sent=0
rcvd=0
shaper_drop_sent=0
shaper_drop_rcvd=0
perip_drop=0
src_int="internal"
dst_int="wan1"
SN=97404 app="N/A"
app_cat="N/A"
carrier_ep="N/A"

```

If you want to know more about logging, see the Logging and Reporting chapter in the FortiOS Handbook. If you want to know more about traffic log messages, see the FortiGate Log Message Reference.

Quality of service

The Quality of Service (QoS) feature allows the management of the level of service and preference given to the various types and sources of traffic going through the firewall so that the traffic that is important to the services and functions connecting through the firewall gets the treatment required to ensure the level of quality that is required.

QoS uses the following techniques:

Traffic policing	Packets are dropped that do not conform to bandwidth limitations
Traffic Shaping	Assigning minimum levels of bandwidth to be allocated to specific traffic flows to guarantee levels of servers or assigning maximum levels of bandwidth to be allocated to specific traffic flows so that they do not impede other flows of traffic.

This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.

Queuing

Assigning differing levels priority to different traffic flows so that traffic flows that are adversely effected by latency are prevented from being effected by traffic flows that are not subject to the effects of latency. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

An example of where you would want to use something like this is if you had competing traffic flows of Voice over IP traffic and email traffic. The VoIP traffic is highly susceptible to latency issues. If you have a delay of a few seconds it is quickly noticeable when it is occurring. Email on the other hand can have a time delay of much longer and it is highly unlikely that it will be noticed at all.



By default, the priority given to any traffic is high, so if you want to give one type of traffic priority over all other traffic you will need to lower the priority of all of the other traffic.

Policy Monitor

Once policies have been configured and enabled it is useful to be able to monitor them. To get an overview about what sort of traffic the policies are processing go to Policy > Monitor > Policy Monitor.

The window is separated into two panes.

Upper pane

The upper pane displays a horizontal bar graph comparing the **Top Policy Usage** based on one of the following criteria:

- Active Sessions
- Bytes
- Packets

The criteria that the displayed graph is based on can be selected from the drop down menu in the upper right corner of the pane. The field name is **Report By:**.

The bars of the graph are interactive to an extent and can be used to drill down for more specific information. If you hover the cursor over the bar of the graph a small popup box will appear displaying more detailed information. If the bar of the graph is selected an entirely new window will be displayed using a vertical bar graph to divide the data that made up the first graph by IP address.

For example if the first graph was reporting usage by active sessions it would include a bar for each of the top policies with a number at the end showing how many sessions were currently going through that policy. If one of the bars of the graph was then selected the new bar graph would show the traffic of that policy separated by either **Source Address**, **Destination Address** or **Destination Port**. As in the other window, the selection for the reported criteria is in the upper right corner of the pane. If the parameter was by source address there would be a bar for each of the IP addresses sending a session through the policy and the end of the bar would show how many sessions.

To go back to the previous window of information in the graphs select the Return link in the upper left of the pane.

Lower pane

The lower pane contains a spreadsheet of the information that the bar graph will derive their information from. The column headings will include:

- Policy ID
- Source Interface/Zone

- Destination Interface/Zone
- Action
- Active Sessions
- Bytes
- Packets

Network defense

This section describes in general terms the means by which attackers can attempt to compromise your network and steps you can take to protect it. The goal of an attack can be as complex as gaining access to your network and the privileged information it contains, or as simple as preventing customers from accessing your web server. Even allowing a virus onto your network can cause damage, so you need to protect against viruses and malware even if they are not specifically targeted at your network.

The following topics are included in this section:

- Monitoring
- Blocking external probes
- Defending against DoS attacks

Monitoring

Monitoring, in the form of logging, alert email, and SNMP, does not directly protect your network. But monitoring allows you to review the progress of an attack, whether afterwards or while in progress. How the attack unfolds may reveal weaknesses in your preparations. The packet archive and sniffer policy logs can reveal more details about the attack. Depending on the detail in your logs, you may be able to determine the attackers location and identity.

While log information is valuable, you must balance the log information with the resources required to collect and store it.

Blocking external probes

Protection against attacks is important, but attackers often use vulnerabilities and network tools to gather information about your network to plan an attack. It is often easier to prevent an attacker from learning important details about your network than to defend against an attack designed to exploit your particular network.

Attacks are often tailored to the hardware or operating system of the target, so reconnaissance is often the first step. The IP addresses of the hosts, the open ports, and the operating systems the hosts are running is invaluable information to an attacker. Probing your network can be as simple as an attacker performing an address sweep or port scan to a more involved operation like sending TCP packets with invalid combinations of flags to see how your firewall reacts.

Address sweeps

An address sweep is a basic network scanning technique to determine which addresses in an address range have active hosts. A typical address sweep involves sending an ICMP ECHO request (a ping) to each address in an address range to attempt to get a response. A response signifies that there is a host at this address that responded to the ping. It then becomes a target for more detailed and potentially invasive attacks.

Address sweeps do not always reveal all the hosts in an address range because some systems may be configured to ignore ECHO requests and not respond, and some firewalls and gateways may be configured to prevent ECHO requests from being transmitted to the destination network. Despite this shortcoming, Address sweeps are still used because they are simple to perform with software tools that automate the process.

Use the `icmp_sweep` anomaly in a DoS policy to protect against address sweeps.

There are a number of IPS signatures to detect the use of ICMP probes that can gather information about your network. These signatures include `AddressMask`, `Traceroute`, `ICMP.Invalid.Packet.Size`, and `ICMP.Oversized.Packet`. Include ICMP protocol signatures in your IPS sensors to protect against these probes/attacks.

Port scans

Potential attackers may run a port scan on one or more of your hosts. This involves trying to establish a communication session to each port on a host. If the connection is successful, a service may be available that the attacker can exploit.

Use the DoS anomaly check for `tcp_port_scan` to limit the number of sessions (complete and incomplete) from a single source IP address to the configured threshold. If the number of sessions exceed the threshold, the configured action is taken.

Use the DoS anomaly check for `udp_scan` to limit UDP sessions in the same way.

Probes using IP traffic options

Every TCP packet has space reserved for eight flags or control bits. They are used for communicating various control messages. Although space in the packet is reserved for all eight, there are various combinations of flags that should never happen in normal network operation. For example, the SYN flag, used to initiate a session, and the FIN flag, used to end a session, should never be set in the same packet.

Attackers may create packets with these invalid combinations to test how a host will react. Various operating systems and hardware react in different ways, giving a potential attackers clues about the components of your network.

The IPS signature `TCP.Bad.Flags` detects these invalid combinations. The default action is pass though you can override the default and set it to **Block** in your IPS sensor.

Configure packet replay and TCP sequence checking

The anti-replay CLI command allows you to set the level of checking for packet replay and TCP sequence checking (or TCP Sequence (SEQ) number checking). All TCP packets contain a Sequence Number (SEQ) and an Acknowledgement Number (ACK). The TCP protocol uses these numbers for error free end-to-end communications. TCP sequence checking can also be used to validate individual packets.

FortiGate units use TCP sequence checking to make sure that a packet is part of a TCP session. By default, if a packet is received with sequence numbers that fall out of the expected range, the FortiGate unit drops the packet. This is normally a desired behavior, since it means that the packet is invalid. But in some cases you may want to configure different levels of anti-replay checking if some of your network equipment uses non-RFC methods when sending packets.

Configure the anti-replay CLI command:

```
config system global
    set anti-replay {disable | loose | strict}
```

end

You can set anti-replay protection to the following settings:

- `disable` — No anti-replay protection.
- `loose` — Perform packet sequence checking and ICMP anti-replay checking with the following criteria:
 - The SYN, FIN, and RST bit can not appear in the same packet.
 - The FortiGate unit does not allow more than one ICMP error packet through before it receives a normal TCP or UDP packet.
 - If the FortiGate unit receives an RST packet, and `check-reset-range` is set to `strict`, the FortiGate unit checks to determine if its sequence number in the RST is within the un-ACKed data and drops the packet if the sequence number is incorrect.
- `strict` — Performs all of the loose checking but for each new session also checks to determine if the TCP sequence number in a SYN packet has been calculated correctly and started from the correct value for each new session. Strict anti-replay checking can also help prevent SYN flooding.

If any packet fails a check it is dropped.

Configure ICMP error message verification

Enable ICMP error message verification to ensure an attacker can not send an invalid ICMP error message.

```
config system global
    check-reset-range {disable | strict}
end
```

- `disable` — the FortiGate unit does not validate ICMP error messages.
- `strict` — enable ICMP error message checking.

If the FortiGate unit receives an ICMP error packet that contains an embedded IP(A,B) | TCP(C,D) header, then if FortiOS can locate the A:C->B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped. Strict checking also affects how the anti-replay option checks packets.

Protocol header checking

Select the level of checking performed on protocol headers.

```
config system global
    check-protocol-header {loose | strict}
end
```

- `loose` — the FortiGate unit performs basic header checking to verify that a packet is part of a session and should be processed. Basic header checking includes verifying that the layer-4 protocol header length, the IP header length, the IP version, the IP checksum, IP options are correct, etc.
- `strict` — the FortiGate unit does the same checking as above plus it verifies that ESP packets have the correct sequence number, SPI, and data length.

If the packet fails header checking it is dropped by the FortiGate unit.

Evasion techniques

Attackers employ a wide range of tactics to try to disguise their techniques. If an attacker disguises a known attack in such a way that it is not recognized, the attack will evade your security and possibly succeed. FortiGate security recognizes a wide variety of evasion techniques and normalizes data traffic before inspecting it.

Packet fragmentation

Information sent across local networks and the Internet is encapsulated in packets. There is a maximum allowable size for packets and this maximum size varies depending on network configuration and equipment limitations. If a packet arrives at a switch or gateway and it is too large, the data it carries is divided among two or more smaller packets before being forwarded. This is called fragmentation.

When fragmented packets arrive at their destination, they are reassembled and read. If the fragments do not arrive together, they must be held until all of the fragments arrive. Reassembly of a packet requires all of the fragments.

The FortiGate unit automatically reassembles fragmented packets before processing them because fragmented packets can evade security measures. Both IP packets and TCP packets are reassembled by the IPS engine before examination.

For example, you have configured the FortiGate unit to block access to the example.org web site. Any checks for example.com will fail if a fragmented packet arrives and one fragment contains `http://www.exa` while the other contains `mple.com/`. Viruses and malware can be fragmented and avoid detection in the same way. The FortiGate unit will reassemble fragmented packets before examining network data to ensure that inadvertent or deliberate packet fragmentation does not hide threats in network traffic.

Non-standard ports

Most traffic is sent on a standard port based on the traffic type. The FortiGate unit recognizes most traffic by packet content rather than the TCP/UDP port and uses the proper IPS signatures to examine it. Protocols recognized regardless of port include DHCP, DNP3, FTP, HTTP, IMAP, MS RPC, NNTP, POP3, RSTP, SIP, SMTP, and SSL, as well as the supported IM/P2P application protocols.

In this way, the FortiGate unit will recognize HTTP traffic being sent on port 25 as HTTP rather than SMTP, for example. Because the protocol is correctly identified, the FortiGate unit will examine the traffic for any enabled HTTP signatures.

Negotiation codes

Telnet and FTP servers and clients support the use of negotiation information to allow the server to report what features it supports. This information has been used to exploit vulnerable servers. To avoid this problem, the FortiGate unit removes negotiation codes before IPS inspection.

HTTP URL obfuscation

Attackers encode HTML links using various formats to evade detection and bypass security measures. For example, the URL `www.example.com/cgi.bin` could be encoded in a number of ways to avoid detection but still work properly, and be interpreted the same, in a web browser.

The FortiGate prevents the obfuscation by converting the URL to ASCII before inspection.

HTTP URL obfuscation types

Encoding type	Example
No encoding	<code>http://www.example.com/cgi.bin/</code>

Encoding type	Example
Decimal encoding	http://www.example.com/c g i . b i n /
URL encoding	http://www.example.com/%43%47%49%2E%42%49%4E%2F
ANSI encoding	http://www.example.com/%u0063%u0067%u0069%u002E%u0062%u0069%u006E/
Directory traversal	http://www.example.com/cgi.bin/test/..

HTTP header obfuscation

The headers of HTTP requests or responses can be modified to make the discovery of patterns and attacks more difficult. To prevent this, the FortiGate unit will:

- remove junk header lines
- reassemble an HTTP header that's been folded onto multiple lines
- move request parameters to HTTP POST body from the URL

The message is scanned for any enabled HTTP IPS signatures once these problems are corrected.

HTTP body obfuscation

The body content of HTTP traffic can be hidden in an attempt to circumvent security scanning. HTTP content can be GZipped or deflated to prevent security inspection. The FortiGate unit will uncompress the traffic before inspecting it.

Another way to hide the contents of HTTP traffic is to send the HTTP body in small pieces, splitting signature matches across two separate pieces of the HTTP body. The FortiGate unit reassembles these 'chunked bodies' before inspection.

Microsoft RPC evasion

Because of its complexity, the Microsoft Remote Procedure Call protocol suite is subject to a number of known evasion techniques, including:

- SMB-level fragmentation
- DCERPC-level fragmentation
- DCERPC multi-part fragmentation
- DCERPC UDP fragmentation
- Multiple DCERPC fragments in one packet

The FortiGate unit reassembles the fragments into their original form before inspection.

Defending against DoS attacks

A denial of service is the result of an attacker sending an abnormally large amount of network traffic to a target system. Having to deal with the traffic flood slows down or disables the target system so that legitimate users can not use it for the duration of the attack.

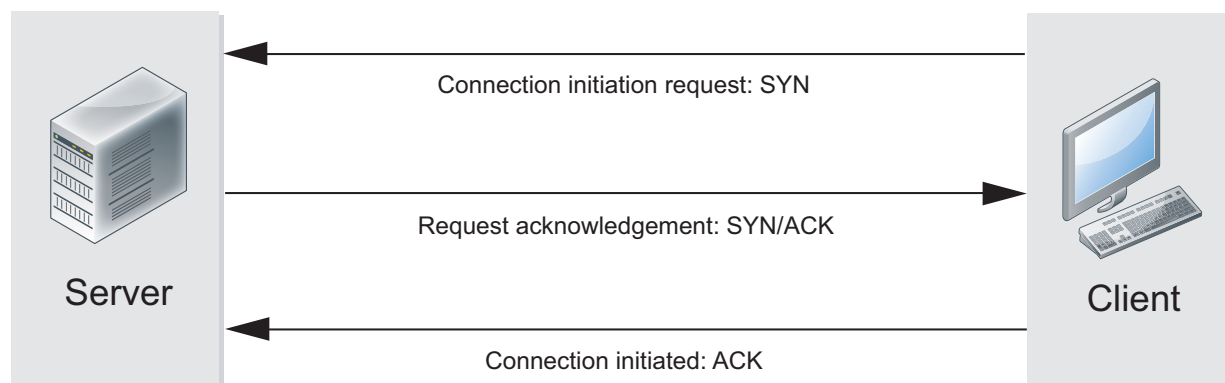
Any network traffic the target system receives has to be examined, and then accepted or rejected. TCP, UDP, and ICMP traffic is most commonly used, but a particular type of TCP traffic is the most effective. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are started between systems.

The “three-way handshake”

Communication sessions between systems start with establishing a TCP/IP connection. This is a simple three step process, sometimes called a “three-way handshake,” initiated by the client attempting to open the connection.

1. The client sends a TCP packet with the SYN flag set. With the SYN packet, the client informs the server of its intention to establish a connection.
2. If the server is able to accept the connection to the client, it sends a packet with the SYN and the ACK flags set. This simultaneously acknowledges the SYN packet the server has received, and informs the client that the server intends to establish a connection.
3. To acknowledge receipt of the packet and establish the connection, the client sends an ACK packet.

Establishing a TCP/IP connection



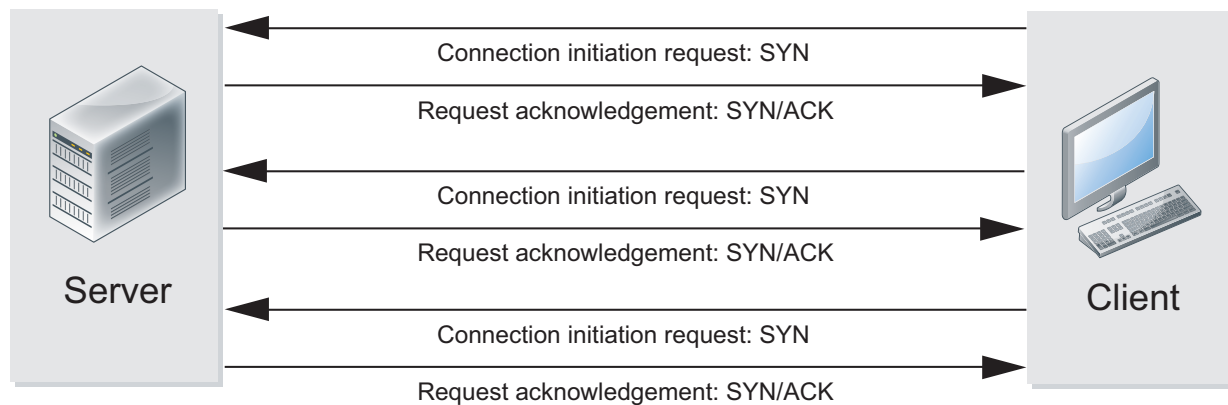
The three-way handshake is a simple way for the server and client to each agree to establish a connection and acknowledge the other party expressing its intent. Unfortunately, the three-way handshake can be used to interfere with communication rather than facilitate it.

SYN flood

When a client sends a SYN packet to a server, the server creates an entry in its session table to keep track of the connection. The server then sends a SYN+ACK packet expecting an ACK reply and the establishment of a connection.

An attacker intending to disrupt a server with a denial of service (DoS) attack can send a flood of SYN packets and not respond to the SYN+ACK packets the server sends in response. Networks can be slow and packets can get lost so the server will continue to send SYN+ACK packets until it gives up, and removes the failed session from the session table. If an attacker sends enough SYN packets to the server, the session table will fill completely, and further connection attempts will be denied until the incomplete sessions time out. Until this happens, the server is unavailable to service legitimate connection requests.

A single client launches a SYN flood attack

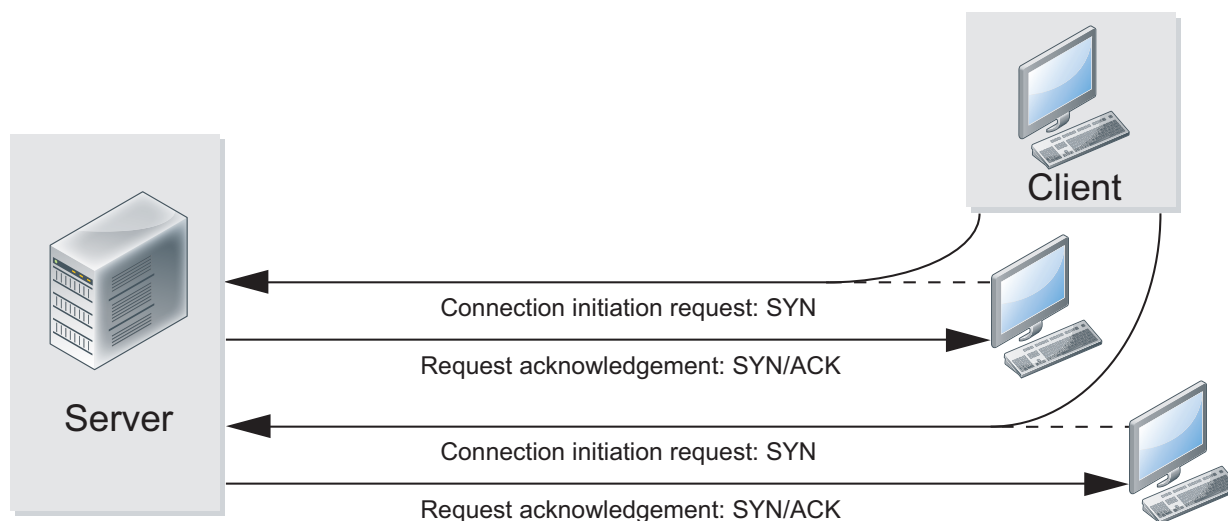


SYN floods are seldom launched from a single address so limiting the number of connection attempts from a single IP address is not usually effective.

SYN spoofing

With a flood of SYN packets coming from a single attacker, you can limit the number of connection attempts from the source IP address or block the attacker entirely. To prevent this simple defense from working, or to disguise the source of the attack, the attacker may spoof the source address and use a number of IP addresses to give the appearance of a distributed denial of service (DDoS) attack. When the server receives the spoofed SYN packets, the SYN+ACK replies will go to the spoofed source IP addresses which will either be invalid, or the system receiving the reply will not know what to do with it.

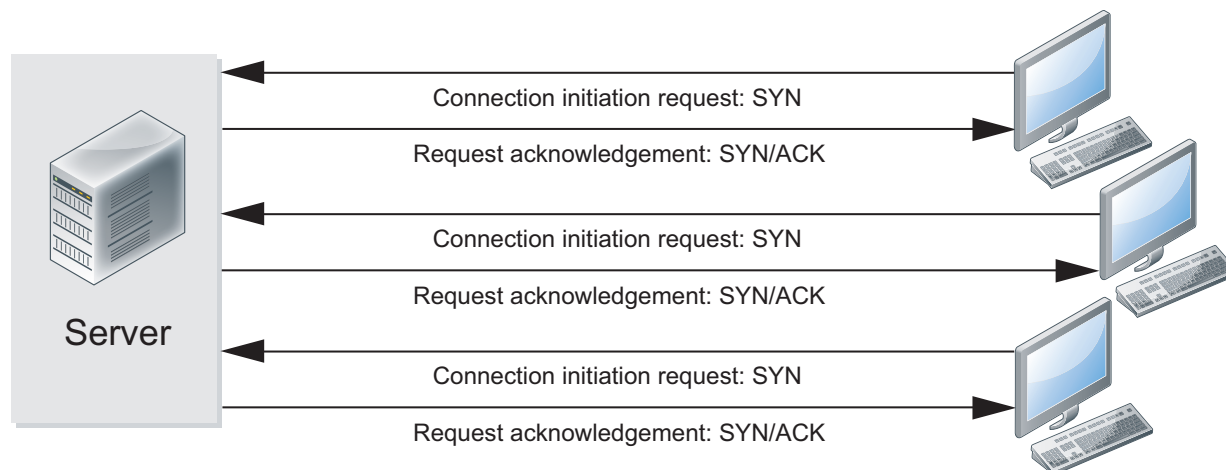
A client launches a SYN spoof attack



DDoS SYN flood

The most severe form of SYN attack is the distributed SYN flood, one variety of distributed denial of service attack (DDoS). Like the SYN flood, the target receives a flood of SYN packets and the ACK+SYN replies are never answered. The attack is distributed across multiple sources sending SYN packets in a coordinated attack.

Multiple attackers launch a distributed SYN flood



The distributed SYN flood is more difficult to defend against because multiple clients are capable of creating a larger volume of SYN packets than a single client. Even if the server can cope, the volume of traffic may overwhelm a point in the network upstream of the targeted server. The only defence against this is more bandwidth to prevent any choke-points.

Configuring the SYN threshold to prevent SYN floods

The preferred primary defence against any type of SYN flood is the DoS anomaly check for `tcp_syn_flood` threshold. The threshold value sets an upper limit on the number of new incomplete TCP connections allowed per second. If the number of incomplete connections exceeds the threshold value, and the action is set to **Pass**, the FortiGate unit will allow the SYN packets that exceed the threshold. If the action is set to **Block**, the FortiGate unit will block the SYN packets that exceed the threshold, but it will allow SYN packets from clients that send another SYN packet.

The tools attackers use to generate network traffic will not send a second SYN packet when a SYN+ACK response is not received from the server. These tools will not “retry.” Legitimate clients will retry when no response is received, and these retries are allowed even if they exceed the threshold with the action set to **Block**.

SYN proxy

FortiGate units with network acceleration hardware, whether built-in or installed in the form of an add-on module, offer a third action for the `tcp_syn_flood` threshold. Instead of **Block** and **Pass**, you can choose to **Proxy** the incomplete connections that exceed the threshold value.

When the `tcp_syn_flood` threshold action is set to **f**, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the FortiGate unit will intercept incoming SYN packets from clients and respond with a SYN+ACK packet. If the FortiGate unit receives an ACK

response as expected, it will “replay” this exchange to the server to establish a communication session between the client and the server, and allow the communication to proceed.

Other flood types

UDP and ICMP packets can also be used for DoS attacks, though they are less common. TCP SYN packets are so effective because the target receives them and maintains a session table entry for each until they time out. Attacks using UDP or ICMP packets do not require the same level of attention from a target, rendering them less effective. The target will usually drop the offending packets immediately, closing the session.

Use the `udp_flood` and `icmp_flood` thresholds to defend against these DoS attacks.

DoS policies

DDoS attacks vary in nature and intensity. Attacks aimed at saturating the available bandwidth upstream of your service can only be countered by adding more bandwidth. DoS policies can help protect against DDoS attacks that aim to overwhelm your server resources.

DoS policy recommendations

- Use and configure DoS policies to appropriate levels based on your network traffic and topology. This will help drop traffic if an abnormal amount is received.
- It is important to set a good threshold. The threshold defines the maximum number of sessions/packets per second of normal traffic. If the threshold is exceeded, the action is triggered. Threshold defaults are general recommendations, although your network may require very different values.
- One way to find the correct values for your environment is to set the action to **Pass** and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the more protected your system will be from DoS attacks, but your system will also be more likely to generate false alarms.

GUI and CLI - what you may not know

The Graphic User Interface (GUI) is designed to be as intuitive as possible but there are always a few things that are left out because to put all of that information on the interface would clutter it up to the point where it wouldn't be graphical and intuitive anymore.

This section is made up of knowledge that will make working with the both of the management interfaces easier because you won't have to find out about things like field limitations through trial and error. Some of it has to do with changing in how navigation in the GUI has changed.

The section includes the topics:

- Mouse Tricks
- Changing the default column setting on the policy page
- Naming Rules and Restrictions
- Character Restrictions
- Length of Fields Restrictions
- Object Tagging and Coloring
- Numeric Values
- Selecting options from a list
- Enabling or disabling options
- To Enable or Disable Optionally Displayed Features

Mouse tricks

In previous version of the firmware much of the navigation, editing or choosing of options in the Web-based Manager was carried out by using the mouse in combination with a number of icons visible on the interface. This version of the firmware makes more extensive use of the right or secondary mouse button as well as the “drag and drop” feature. If you are used to the old Web-based Manager interface you will notice that a number of the options at the top of the display window are not there anymore or there are fewer of them.

To get a feel for the new approach the **Policy & Objects > Policy > IPv4** window is a noticeable place to see some of these changes in action.

The different view modes are still in the upper right-hand corner as they were before but now there is no column settings link to move or configure the columns of the window. Now if you wish to reposition a column just use the mouse to click on the column heading and drag it to its new position. If you wish to add a new column just right-click on one of the column headings and a drop down menu will appear with the option “Column Settings”. Use the right pointing triangle to expand the “Column Settings” option to see a choice of possible columns for the window you are in. Those already selected will be at the top with a checked box and the available new ones will be at the bottom ready to be selected.

Rather than having a link to initiate a move in the positioning of policies in the sequence, you can select a policy and hold down the mouse button and drag it to its new position.

By right or secondary clicking the mouse curser in the cells of the Policy window you will get a drop down menu that is contextual to the column and policy row where you made the click. For example if you right click in the “Schedule” column for the row that is for policy #5 you will get the option to select a schedule for policy #5 along with a number of other configuration options relating to that policy or its position in the sequence of policies.

You will find this approach used much more frequently through out the Web-based Manager, giving it a more modern and intuitive feel once you learn to use the right mouse button rather than finding a link displayed on the page.

Changing the default column setting on the policy page

The **Policy & Objects > Policy > IPv4** window is one of the more important ones in the Web based interface and has the capacity to display a lot of information, but displaying all of that information at the same time makes for a very busy screen. If all of the columns are displayed, depending on the screen size you may have to constantly use the scroll bars to see what you need to look at. The default installation shows some of the more commonly used columns but these list may not consist of the columns that you wish to look at or the order that you wish to view them in. For this reason it is possible, through the CLI to override these settings to establish a new default.

The syntax of the command starts with:

```
config system settings
set gui-default-policy-columns
```

The rest of the command is a space delimited list that depends on the columns you wish to view and the order you wish to view them in. The possible selection is in the following table.

Variables for the gui-default-policy-columns command

Variable Name	Column Heading
#	Sequence Number
policyid	Policy ID
srcintf	Source Interface
dstintf	Destination Interface
srcaddr	Source Addresses
dstaddr	Destination Addresses
schedule	Policy Schedule
service	Policy Services
action	Policy Action
logtraffic	Traffic Logging Status
nat	Policy NAT Status
status	Policy Status
authentication	Authentication Groups

Variable Name	Column Heading
count	Policy Traffic Counter
profile	Security Profiles
vpntunnel	VPN Tunnel
comments	Policy Comment

Example

If you wanted these columns in this order, Policy ID, Source Addresses, Destination Addresses, Security Profiles, Policy Comment. You would enter the command:

```
config system settings
  set gui-default-policy-columns policyid srcaddr dstaddr profile comments
```

Naming rules and restrictions

The following are the specific rules that are obeyed by the FortiGate.

Duplicate Name Issues:

- A VLAN cannot have the same name as a physical interface.
- An Address must not have the same name as an Address Group.
- An Address or Address Group must not have the same name as a Virtual IP Address.
- A Service cannot have the same name as a Service Group.
- A VLAN must not have the same name as a VDOM.
- A VLAN or VDOM must not have the same name as a Zone.



Try to make each firewall object name as unique as possible so that it cannot be confused with another object.

Character restrictions

Object names cannot contain the following characters:

- <
- >
- (
-)

All other characters, including other languages, are supported as usable characters.



FortiOS allows spaces in just about all object name fields, but this was not always the case. Some people are still cautious where this is concerned. If you're cautious about characters in the names of objects use this basic rule of thumb:

When naming objects, only use characters that are alphanumeric (a-z, A-Z, 0-9) and where there is the temptation to use spaces in a name, use the '-' (dash) and '_' (underscore).

Numeric values

Numeric values are used to configure various sizes, rates, numeric addresses, or other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or as in the case of MAC or IPv6 addresses separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

Numeric values

Numeric values are used to configure various sizes, rates, numeric addresses, or other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or as in the case of MAC or IPv6 addresses separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

Selecting options from a list

If a configuration field can only contain one of a number of selected options, the web-based manager and CLI present you a list of acceptable options and you can select one from the list. No other input is allowed. From the CLI you must spell the selection name correctly.

Enabling or disabling options

If a configuration field can only be on or off (enabled or disabled) the web-based manager presents a check box or other control that can only be enabled or disabled. From the CLI you can set the option to enable or disable.

To enable or disable optionally displayed features

There are a number of features in the web-based manager that can be configured to either be displayed if you are likely to use them or disabled if you have no need to see them. The ones that may be relevant to the function of the Firewall are:

- Central NAT Table
- Dynamic Profile
- Explicit Proxy
- Implicit Firewall Policies
- IPv6
- Load Balance
- Local In Policy

You can enable or disable these features by going to **System > Admin > Settings** or by using the following CLI options:

```
config system global
  set gui-ap-profile {disable | enable}
  set gui-central-nat-table {disable | enable}
  set gui-dns-database {disable | enable}
  set gui-dynamic-profile-display {disable | enable}
  set gui-icap {disable | enable}
  set gui-implicit-id-based-policy {disable | enable}
  set gui-implicit-policy {disable | enable}
  set gui-ipsec-manual-key {enable | disable}
  set gui-ipv6 {enable | disable}
  set gui-lines-per-page <gui_lines>
  set gui-load-balance {disable | enable}
  set gui-object-tags {disable | enable}
  set gui-policy-interface-pairs-view {enable | disable}
  set gui-voip-profile {disable | enable}
end
```

Building firewall objects and policies

The other chapters in the Firewall book have so far been concerned primarily with concepts and abstract ideas that are designed to help you understand what is going on with the firewall and what it can do. Now that we have a good grounding in the “what” it is time to get into the “how”.

This section will provide the instructions for the web-based manager (when available) and the CLI for adding and or editing FortiGate firewall objects and then how to put them together when building a policy to govern the traffic flowing through your network. To give some context, scenarios have been included. The instructions here are concerned with the creation of the objects. The inclusion of these objects into firewall policies is not shown in these instructions.

This chapter includes the instructions for building the following:

- IPv4 Firewall Addresses
- IPv6 Firewall Addresses
- FQDN address
- Changing the TTL of a FQDN address
- New Geography-based Address
- IPv4 Address Group
- IPv4 Address Group
- IPv6 Address Group
- Multicast Address
- Service Category
- TCP/UDP/SCTP Service
- ICMP Service
- ICMPv6 Service
- Service Group
- Virtual IP address
- IP Pool
- Central NAT Table
- Firewall Schedule - Recurring
- Firewall Schedule - One-time
- Schedule Group
- Proxy Option
- DoS Policy

Building firewall objects and policies

The other chapters in the Firewall book have so far been concerned primarily with concepts and abstract ideas that are designed to help you understand what is going on with the firewall and what it can do. Now that we have a good grounding in the “what” it is time to get into the “how”.

This section will provide the instructions for the web-based manager (when available) and the CLI for adding and/or editing FortiGate firewall objects and then how to put them together when building a policy to govern the traffic flowing through your network. To give some context, scenarios have been included. The instructions here are concerned with the creation of the objects. The inclusion of these objects into firewall policies is not shown in these instructions.

This chapter includes the instructions for building the following:

- IPv4 Firewall Addresses
- IPv6 Firewall Addresses
- FQDN address
- Changing the TTL of a FQDN address
- New Geography-based Address
- IPv4 Address Group
- IPv4 Address Group
- IPv6 Address Group
- Multicast Address
- Service Category
- TCP/UDP/SCTP Service
- ICMP Service
- ICMPv6 Service
- Service Group
- Virtual IP address
- IP Pool
- Central NAT Table
- Firewall Schedule - Recurring
- Firewall Schedule - One-time
- Schedule Group
- Proxy Option
- DoS Policy

Example: IPv4 firewall addresses

Scenario: Mail server

You need to create an IPv4 address for the Mail Server on Port1 of your internal network.

- These server is on the network off of port1.
 - The IP address is 192.168.1.27
 - The subnet mask is 255.255.255.0
 - There should be a tag for this address being for a server
1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
 2. Fill out the fields with the following information:

Category	Address
Name	Mail_Server
Type	Subnet / IP Range
Subnet / IP Range	192.168.1.27/255.255.255.0
Interface	port1
Visibility	<enabled>
Comments	<add any additional discription>

3. Select **OK**.
4. Enter the following CLI command:

```
config firewall address
edit Mail_Server
set type ipmask
set subnet 192.168.1.27 255.255.255.255
set associated-interface port1
end
```

Scenario: First floor network

You need to create an IPv4 address for the subnet of the internal network off of Port1.

- These computers are on the network off of port1.
 - The subnet is the range from 192.168.1.1 to 192.168.1.255.
 - The subnet mask is 255.255.255.0
 - There should be a reference to this being the network for the 1st floor of the building.
1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
 2. Fill out the fields with the following information

Category	Address
Name	Internal_Subnet_1
Type	Subnet
Subnet / IP Range	192.168.1.0/24
Interface	port1

visibility	<enabled>
Comments	Network for 1st Floor

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall address
edit Internal_Subnet_1
    set comment "Network for 1st Floor"
    set type ipmask
    set subnet 192.168.1.0/24
    set associated-interface port1
end
```

Scenario: Marketing department

You need to create an IPv4 address for the address range for a group of computers used by the Marketing Department.

- These computers are on the network off of port1.
- The IP addresses for these computers range from 192.168.1.100 to 192.168.1.115

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information

Category	Address
Name	Marketing_computers
Type	IP Range
Subnet / IP Range	192.168.1.100-192.168.1.115
Interface	port1
Visibility	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall address
edit Internal_Subnet_1
    set type iprange
    set start-ip 192.168.1.100
    set end-ip 192.168.1.115
    set associated-interface port1
end
```

Verification

To verify that the addresses were added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:


```
config firewall address
  edit <the name of the address to verify>
  show full-configuration
```

Example: IPv6 firewall addresses

Scenario: Mail server

You need to create an IPv6 address for the Mail Server on Port1 of your internal network.

- These server is on the network off of port1.
 - The IP address is 2001:db8:0:2::20/64
 - There should be a tag for this address being for a server
1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
 2. Fill out the fields with the following information

Category	IPv6 Address
Name	Mail_Server
Type	Subnet
Subnet / IP Range	2001:db8:0:2::20/64
Interface	port1
Visibility	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.
4. Enter the following CLI command:


```
config firewall address6
  edit Mail_Server
    set type ipmask
    set subnet 2001:db8:0:2::20/64
    set associated-interface port1
  end
```

Scenario: First floor network

You need to create an IPv4 address for the subnet of the internal network off of Port1.

- These computers are on the network off of port1.
- The Network uses the IPv6 addresses: fdde:5a7d:f40b:2e9d:xxxx:xxxx:xxxx:xxxx
- There should be a reference to this being the network for the 1st floor of the building.

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information

Category	IPv6 Address
Name	Internal_Subnet_1
Type	Subnet
Subnet / IP Range	2001:db8:0:2::/64
Interface	port1
Visibility	<enable>
Comments	Network for 1st Floor

3. Select **OK**.
4. Enter the following CLI command:


```
config firewall address6
edit Internal_Subnet_1
    Set comment "Network for 1st Floor"
    set subnet 2001:db8:0:2::/64
end
```

Verification

To verify that the addresses were added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:


```
config firewall address6
edit <the name of the address that you wish to verify>
    Show full-configuration
```

Example: FQDN address

You have to great a policy that will govern traffic that goes to a site that has a number of servers on the Internet. Depending on the traffic or the possibility that one of the servers is down network traffic can go to any one of those sites. The consistent factor is that they all use the same Fully Qualified Domain Name.

- The FQDN of the web site: example.com
 - The number of ISP connections off of the FortiGate firewall: 2
1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
 2. Fill out the fields with the following information:

Category	Address
-----------------	---------

Name	BigWebsite.com
Type	FQDN
FQDN	bigwebsite.com
Interface	any
Visibility	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall address
edit BigWebsite.com
set type fqdn
set associated-interface any
set fqdn bigwebsite.com
end
```

Verification

To verify that the addresses were added correctly:

1. Go to **Firewall Objects > Address > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:

```
config firewall address
edit <the name of the address that you wish to verify>
Show full-configuration
```

Changing the TTL of a FQDN address

To make sure that the FQDN resolves to the most recent active server you have been asked to make sure that the FortiGate has not cached the address for any longer than 10 minutes.

There is no field for the cached time-to-live in the web-based manager. It is only configurable in the CLI. Enter the following commands:

```
config firewall address
edit BigWebsite.com
set cache-ttl 600
end
```

Example: Geography-based address

Your company is US based and has information on its web site that may be considered information that is not allowed to be sent to embargoed countries. In an effort to help reduce the possibility of sensitive information

going to those countries you have been asked to set up addresses for those countries so that they can be blocked in the firewall policies.

- One of the countries you have been asked to block is Cuba
- You have been asked to comment the addresses so that other administrators will know why they have been created

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information

Category	Address
Name	Cuba
Type	Geography
Country	Cuba
Interface	any
Visibility	<enable>
Comments	Embargoed

3. Select **OK**.
4. Enter the following CLI command:

```
config firewall address
edit Cuba
set type geography
set country CN
set interface wan1
end
```

Example: IPv4 address group

Your company has a small division that is in Denmark that has a number of remote users that need to connect to a resource from either home, office or customer sites. The thing that they have in common is that there are all in Denmark. An address group needs to be created that will allow for this.

The preconfigured addresses to use will consist of:

- Denmark - a geography based address
- Denmark_ISP1 - a IP range of address of an ISP that services Denmark
- Denmark_ISP2 - a IP range of address of another ISP that services Denmark
- Denmark_Division - the FQDN of the Denmark office that uses Dynamic DNS

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address Group**.
2. Fill out the fields with the following information

Type	IPv4 Group
-------------	------------

Group Name	Denmark_Users
Show in Address List	<enable>
Members	Denmark Denmark_ISP1 Denmark_ISP2 Denmark_Division
Comments	<Input into this field is optional>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall addrgrp
edit Denmark_Users
set member Denmark Denmark_ISP1 Denmark_ISP2 Denmark_Division
end
```

To verify that the addresses were added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.

2. Enter the following CLI command:

```
config firewall addgrp
edit <the name of the address that you wish to verify>
Show full-configuration
```

Example: IPv6 address group

Create IPv6 address groups from existing IPv6 addresses

Your company has 3 internal servers with IPv6 addresses that it would like to group together for the purposes of a number of policies.

The preconfigured addresses to use will consist of:

- Web_Server-1
- Web_Server-2
- Web_Server-3

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address Group**.

2. Fill out the fields with the following information

Type	IPv6 Group
Group Name	Web_Server_Cluster

Show in Address List	<enable>
Members	Web_Server-1
	Web_Server-2
	Web_Server-3
Comments	<Input into this field is optional>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall addrgrp6
edit Web_Server_Cluster
set member Web_Server-1 Web_Server-2 Web_Server-3
end
```

To verify that the addresses were added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.

2. Enter the following CLI command:

```
config firewall addgrp6
edit <the name of the address that you wish to verify>
Show full-configuration
```

Example: Multicast address

The company has a large high tech campus that has monitors in many of its meeting rooms. It is common practice for company wide notifications of importance to be done in a streaming video format with the CEO of the company addressing everyone at once.

The video is High Definition quality so takes up a lot of bandwidth. To minimize the impact on the network the network administrators have set things up to allow the use of multicasting to the monitors for these notifications. Now it has to be set up on the FortiGate firewall to allow the traffic.

- The range being used for the multicast is 239.5.0.0 to 239.5.255.255
- The interface on this FortiGate firewall will be on port 4

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information

Category	Multicast Address
Name	Meeting_Room_Displays
Type	Multicast IP Range
Multicast IP Range	239.5.0.0-239.5.255.255

Interface	port4
Visibility	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall multicast-address
edit "meeting_room_display"
set associated-interface "port9"
set start-ip 239.5.0.0
set end-ip 239.5.255.255
set visibility enable
next
end
```

To verify that the address range was added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.

2. Enter the following CLI command:

```
config firewall multicast-address
edit <the name of the address that you wish to verify>
Show full-configuration
```

Example: Service category

Add a new category to the list of Service Categories

You plan on adding a number of devices such as web cameras that will allow the monitoring of the physical security of your datacenter. A number of non-standard services will have to be created and you would like to keep them grouped together under the heading of "Surveillance"

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Category**.

2. Fill out the fields with the following information

Name	Surveillance
Comments	For DataCenter Surveillance Devices

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall service category
Edit Surveillance
Set comment "For DataCenter Surveillance Devices"
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Select the Category Settings icon . A listing of the categories should be displayed.
2. Enter the following CLI command:

```
config firewall service category
show
```

This should bring up all of the categories. Check to see that the new one is displayed.

Example: TCP/UDP/SCTP service**To create and configure a TCP/UDP/STCP protocol type service.**

You have set up some new web cams at work that send a constant live feed to a security service. Not only do these cameras have a feed that can be sent offsite they can be remotely managed from a browser or an application.

The ports that need to be opened to use all of the features of the web cams are:

- Management by browser - TCP on port 8000
- Real time video feed - UDP on port 4000
- Connection through vendor application - SCTP on port 1600

The IP address of the offsite service is 256.25.56.12 (Not a valid IP address. Used for example only)

- One service will be needed for the incoming connections
- One service will be needed for the outgoing connections

The IT manager would like the service for the outgoing data stream to be tied to the destination of the Surveillance service site so that service can only be used for that one vendor.

To add the incoming service

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Service**.
2. Fill out the fields with the following information

Name	WebCam_Connection-incoming
Comments	<Input into this field is optional>
Service Type	Firewall
Show in Service List	Check in check box
Category	Surveillance
Protocol Type	TCP/UDP/SCTP
IP/FQDN	<Leave blank>

Protocol	Destination Port		Source Port	
	Low	High	Low	High
TCP	8000	8000	1	65535
SCTP	16000	16000	1	65535



The source port range can be left blank as the default is 1 to 65535.

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall service custom
edit WebCam_Connection-incoming
Set protocol TCP/UDP/SCTP
Set tcp-portrange 8000
Set sctp-portrange 16000
Set visibility enable
end
```

To add the outgoing service

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Service**.
2. Fill out the fields with the following information

Name	WebCam_Connection-outgoing
Comments	<Input into this field is optional>
Service Type	Firewall
Color	<Changing this value is optional>
Show in Service List	Check in check box
Category	Surveillance
Protocol Type	TCP/UDP/SCTP
IP/FQDN	256.25.56.12

Protocol	Destination Port		Source Port	
	Low	High	Low	High
TCP	4000	4000	1	65535

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall service custom
edit WebCam_Connection-incoming
Set protocol TCP/UDP/SCTP
Set category Surveillance
Set udp-portrange 4000
Set iprange 256.25.56.12
Set visibility enable
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Check that the services have been added to the service list and that they are correct.

2. Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
Show full-configuration
```

This should bring up all of the details of the service.

Example: ICMP service

The Security Officer would like to block the use of the traceroute utility through the network. The IT manager insists that ping and other ICMP utility must be allowed for the task of diagnosing connectivity, so it is agreed that only traceroute functionality will be blocked.

The ICMP type for traceroute is 30. There is no codes with the type.

Web-based Manager Instructions

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Service**.

2. Fill out the fields with the following information

Name	traceroute
Comments	<Input into this field is optional>
Service Type	Firewall
Show in Service List	Check in check box
Category	Uncategorized
Protocol Type	ICMP
Type	30
Code	<Leave blank>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall service custom
edit traceroute
set protocol ICMP
set icmptype 30
set visibility enable
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Check that the services have been added to the services list and that they are correct.
2. Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
show full-configuration
```

Example: ICMPv6 service

The IT Manager is doing some diagnostics and would like to temporarily block the successful replies of ICMP Node information Responses between 2 IPv6 networks.

The ICMP type for ICMP Node informations responses is 140. The codes for a successful response is 0.

Web-based Manager Instructions

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Service**.
2. Fill out the fields with the following information

Name	diagnostic-test1
Comments	<Input into this field is optional>
Service Type	Firewall
Show in Service List	Check in check box
Category	Uncategorized
Protocol Type	ICMP6
Type	140
Code	0

3. Select **OK**.
4. Enter the following CLI command:

```
config firewall service custom
edit diagnostic-test1
set protocol ICMP6
set icmptype 140
set icmpcode 0
set visibility enable
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Check that the services have been added to the services list and that they are correct.
2. Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
show full-configuration
```

Example: Service group

The company provide email services for a number of different companies. They have a standard list of services that they like to keep open to their customer's email servers, including webmail services. The company prides itself on getting a customer up and going the same day so they use standard templates for everything to make sure nothing is forgotten including the services that are available. The services include:

- IMAP
- IMAPS
- POP3
- POP3S
- SMTP
- SMTPS
- HTTP
- HTTPS
- Email_Admin - a custom service for administration of the servers

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Service Group**.
2. Fill out the fields with the following information:

Group Name	Cust_Email_Serv_Template
Comments	(Optional)
Type	Firewall
Members	(click to add...choose from the drop down) <ul style="list-style-type: none"> • IMAP • IMAPS • POP3 • POP3S • SMTP • SMTPS • HTTP • HTTPS • Email_Admin

3. Select **OK**.
4. Enter the following CLI command:

```

config firewall service group
edit Cust_Email_Serv_Template
set member "IMAP" "IMAPS" "POP3" "POP3S" "SMTP" "SMTPS" "HTTP" "HTTPS" "Email_
Admin"
next
end

```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Check that the service group has been added to the services list and that it is correct.
2. Enter the following CLI command:

```

config firewall service group
edit <the name of the service that you wish to verify>
show full-configuration

```

Example: Virtual IP address

The company has an web server on the internal network that needs to be accessed from the Internet.

- The internal IP address is 192.168.50.37
- The external IP address is 256.85.94.60 (for example use only. Not a valid IP address)
- The external IP address is assigned by ISP "A" on WAN1
- The port that needs to be mapped is 80

1. Go to **Policy & Objects > Objects> Virtual IPs** and select **Create New**.
2. Fill out the fields with the following information.

VIP Type	IPv4
Name	Web1-VIP
Comments	Virtual IP for the Forum Webserver
Interface	wan1
Type	(This field can only be changed in the CLI)
Source Address Filter	<disabled>
External IP Address/Range	256.85.94.60
Mapped IP Address/Range	192.168.50.37
Port Forwarding	<enabled>
Protocol	TCP
External Service Port	80
Map to Port	80

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall vip
edit Web1-VIP
    set comment "Virtual IP for the Forum Webserver"
    set extintf wan1
    set extip 256.85.94.60
    set mappedip 192.168.50.37
    set portforward enable
    set protocol tcp
    set extport 80
    set mapped port 80
end
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Virtual IPs**. Check that the virtual IP address has been added to the list and that it is correct.
2. Enter the following CLI command:

```
config firewall vip
edit <the name of the vip that you wish to verify>
show full-configuration
```

VIP group

The company has only a single external IP address but multiple servers with different functions running on its internal LAN that need to be accessed from the Internet.

- The external IP address of the company on wan1 is 256.34.56.149 (for example use only. Not a valid IP address)
 - The webserver is on the internal LAN on 192.168.100.86
 - The webserver needs to answer on ports 80 443
 - The administration of the FortiGate firewall connects on port 4443 instead of 443
 - There is are also a separate email server, FTP server, and Terminal Server for specialised applications.
 - 2 Virtual IPs have been created to map 256.34.56.149 to 192.168.100.86 on ports 80 and 443. The names are webserver_80 and webserver_443 respectively.
1. Go to **Policy & Objects > Objects > Virtual IPs** and select *Create New > Virtual IP Group*.
 2. Fill out the fields with the following information.

Type	IPv4 VIP Group
Name	WebServer_Grp
Comments	(Optional)
Interface	wan1
Members	webserver_80
	webserver_443

3. Enter the following CLI command:

```
config firewall vipgrp
edit WebServer_Grp
set member "webserver_80" "webserver_443"
next
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Virtual IPs**. Check that the virtual IP address group has been added to the list and that it is correct.
2. Enter the following CLI command:

```
config firewall vipgrp
edit <the name of the vip that you wish to verify>
show full-configuration
```

Example: IP pool

Your company has an application server on the internal network that sends out regular data updates to an offsite service. In order to make the service site more secure, they only accept connections from predefined IP address. If the external IP address of the FortiGate firewall interface were used that would mean that the service would be accepting sessions from just about any user in the network so a separate IP address need so be assigned for the Network Address Translation.

- The external address that will be used is one that has been assigned to the company by the ISP on WAN2
- The address is 256.100.42.129 (for example use only. Not a valid IP address)



The ARP interface cannot be set in the Web-based Manager but as this is the only path that the traffic will be coming from the outside this should not be an issue. The setting has been included in the CLI instructions so that you will now how to set it in a situation where you want the ARP replies to be answered only on a specific interface.

1. Go to **Policy & Objects > Objects > IP Pools**.
2. Fill out the fields with the following information:

IP Pool Type	IPv4 Pool
Name	App_Server1
Comments	Addresses assignment for this server only.
Type	One-to-One
External IP Range	256.100.42.129
ARP Reply	<enabled>

3. Select **OK**.
4. Enter the following CLI command:

```
config firewall ippool
edit App_Server1
set comments "Addresses assignment for this server only."
```

```

set type one-to-one
set arp-reply enable
set arp-intf wan2
set startip 256.100.42.129
set endip 256.100.42.129
end

```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > IP Pools**.
2. Check that the IP Pool has been added to the list of IP Pools and that the listed settings are correct.
3. Enter the following CLI command:

```

config firewall ippool
edit <the name of the IP Pool you wish to verify>
show full-configuration

```

Example: Central NAT table

The company has a server on the Development LAN that needs to communicate with a server at a remote site over the Internet. One of the restrictions on the communications between these systems is that the IP address and source port must be specific.

- The traffic going out on to the Internet must be NATed
- The traffic is coming from a server with the IP address 192.168.150.86
- An address called "app-server" has been created for the address 192.168.150.86 on the port1 interface
- The external interface must be 256.23.45.67
- An address called "app-server-ext" has been created for the address 256.23.45.67 on the wan1 interface
- The originating traffic from the server originates in the port range from 2000 to 3000
- The remote site requires that the source TCP port must be within the 12000 to 13000 range

The original address and Translated Address fields require values that are address names that are listed in the address section of Firewall Objects.

1. Go to **Policy & Objects > Policy > Central NAT > Create New**.
2. Fill out the fields with the following information:

Source Address	app-server
Translated Address	app-server-ext
Original Source Port	2000
Translated Port	12000-13000

3. Select **OK**.
4. Enter the following CLI command:

```

config firewall central-nat
edit 0
set orig-addr app-server
set nat-ippool app-server-ext
set orig-port 2000

```



```

set nat-port 12000-13000
next
end

```

To verify that the table was added correctly:

1. Go to **Policy & Objects > Policy > Central NAT**.
2. Check that the table has been added to the list of Central NAT Tables and that the listed settings are correct.
3. Enter the following CLI command:


```

config firewall central-nat
show full-configuration

```
4. Verify that the listing of tables includes the one that you have just configured, with the correct settings.



When configuring the Central NAT in the GUI you may notice that only those addresses which have been configured to be associated with **any** interface are displayed in the drop down menu for choosing a Source Address and yet the CLI will allow any address to be used, not just those associated with **any** interface. This is because by default the policies in the GUI use a function of cross referencing which addresses are allowed based on which interface is involved in the policy. When combined with the aspect of Central NAT that doesn't restrict to a specific interface. This means the only addresses will be allowed are those associated with the **any** interface. The CLI does not have this cross referencing function which is why the CLI seems less restrictive. However, more care must be taken when using the CLI to make sure that appropriate addresses are used.

Example: Firewall schedule - recurring

The Company wants to allow the use of Facebook by employees, but only during none business hours and the lunch break.

- The business hours are 9:00 p.m. to 6:00 p.m.
- The Lunch break is 12:00 p.m. to 1:00 p.m.
- The plan is to create a schedule to cover the morning business hours and the afternoon business hours and block access to the Facebook web site during that time.

1. Go to **Policy & Objects > Objects > Schedule**.
2. Select **Create New > Schedule**.
3. Fill out the fields with the following information:

Type	Recurring
Name	Morning_Business_Hours
Days	Monday, Tuesday, Wednesday, Thursday, Friday
Start Time	Hour = 9, Minute = 0
Stop Time	Hour = 12, Minute = 0

4. Select **OK**.

5. Enter the following CLI command:

```
config firewall schedule recurring
edit Morning_Business_Hours
set day monday tuesday wednesday thursday friday
set start 09:00
set end 12:00
end
```

6. Create a second new schedule.

Type	Recurring
Name	Morning_Business_Hours
Days	Monday, Tuesday, Wednesday, Thursday, Friday
Start Time	Hour = 13, Minute = 0
Stop Time	Hour = 18, Minute = 0

7. Select **OK**.

8. Enter the following CLI command:

```
config firewall schedule recurring
edit Afternoon_Business_Hours
set day monday tuesday wednesday thursday friday
set start 13:00
set end 18:00
end
```

To verify that the schedule was added correctly:

1. Go to **Policy & Objects > Objects > Schedule**.
2. Check that the schedule with the name you used has been added to the list of recurring schedules and that the listed settings are correct.
3. Enter the following CLI command:

```
config firewall schedule recurring
edit <the name of the schedule you wish to verify>
show full-configuration
```

Example: Firewall schedule - one-time

The company wants to change over their web site image to reference the new year. They have decided to take this opportunity to do some hardware upgrades as well. Their web site is business oriented so they have determined that over New Year's Eve there will be very limited traffic.

- They are going to need a maintenance window of 2 hours bracketing midnight on New Year's Eve.

1. Go to **Policy & Objects > Objects > Schedule**.
2. Select **Create New > Schedule**.
3. Fill out the fields with the following information:

Type	One-time
Name	NewYearsEve_Maintenance
Start Date	2014/12/31 <use the built in calendar>
End Date	2015/01/01 <use the built in calendar>
Start Time	Hour: 23, Minute: 0
Stop Time	Hour: 1Minute: 0
Pre-expiration event log	<disable>

4. Select **OK**.
5. Enter the following CLI command:

```
config firewall schedule onetime
edit maintenance_window
set start 23:00 2012/12/31
set end 01:00 2013/01/01
next
end
```

To verify that the schedule was added correctly:

1. Go to **Policy & Objects > Objects > Schedule**.
2. Check that the schedule with the name you used has been added to the list of recurring schedules and that the listed settings are correct.
3. Enter the following CLI command:

```
config firewall schedule onetime
edit <the name of the schedule you wish to verify>
show full-configuration
group you wish to verify>
show full-configuration
```

Example: Schedule group

In order to make the administration of the policies easier a group needs to be created.

1. Go to **Policy & Objects > Objects > Schedule**.
2. Select **Create New > Schedule Group**.
3. Fill out the fields with the following information:

Name	Business_Hours
Members	Morning_Business_Hours Afternoon_Business_Hours

4. Select **OK**.
5. Enter the following CLI command:

```
config firewall service group
edit Business_Hours
```

```

    set member Morning_Business_Hours Afternoon_Business _ours
end

```

To verify that the schedule was added correctly:

1. Go to **Policy & Objects > Objects > Schedule**.
2. Check that the schedule group with the name you used has been added to the list of schedule groups.
3. Enter the following CLI command:

```

config firewall service group
edit <the name of the schedule

```

Example: Proxy option

The company will be using a number of the Security Profiles features on various policies but wants to use as few profiles as possible to make administration simpler. The decision has been made to have two profiles, the default one and a single customized one that will be a combination of the settings required to cover the situations that will not be covered by the default profile.

The company profile will have the following parameters:

- There are no FTP servers running on the site so there is no need for FTP.
- The company has a non-standard IMAP implementation the uses port 1143.
- The Comfort Clients is to be used with a ratio of 1 byte for every 15 seconds.
- There is a lot of varied email traffic so there is to be no blocking of emails due to size beyond the settings on the mail servers.

1. Go to **Policy & Objects > Policy > Proxy Options**.
2. Create a new profile
3. Fill out the fields with the following information:

Name	example_standard
Comments	<optional>
Log Oversized Files	<disabled>

Protocol port mapping:

Enable	Protocol	Inspection Ports
enabled	HTTP	Specify and <leave on default setting.>
enabled	SMTP	Specify and <leave on default setting.>
enabled	POP3	Specify and <leave on default setting.>
enabled	IMAP	Specify and 1143

Enable	Protocol	Inspection Ports
not enabled	FTP	
enabled	NNTP	Specify and <leave on default setting.>
enabled	MAPI	<leave on default setting.>
enabled	DNS	<leave on default setting.>

Common options

Comfort Clients	enabled
Interval (Seconds)	15
Amount(bytes)	1
Block Oversized File/Email	not enabled
Threshold(MB)	not enabled

Web options

Enabled Chunked Bypass	not enabled
Add Fortinet Bar	not enabled
Communication Port	<Unseen because Add Fortinet Bar is not enabled>

Email Options

Allow Fragmented Messages	<not enabled>
Append Signature (SMTP)	<not enabled>
Email Signature Text	<Unseen because Append Signature is not enabled>

4. Select **OK**.

5. Enter the following CLI command:

```
config firewall profile-protocol-options
edit example_standard
config http
set options clientcomfort no-content-summary
set comfort-interval 15
next
config ftp
set status disable
```

```

        set options clientcomfort no-content-summary splice
        set comfort-interval 15
    next
    config imap
        set ports "1143"
        set options fragmail no-content-summary
    next
    config mapi
        set options fragmail no-content-summary
    next
    config pop3
        set options fragmail no-content-summary
    next
    config smtp
        set options fragmail no-content-summary splice
    next
    config nntp
        set options no-content-summary splice
    next
    config ssh
        set inspect-all enable
        set log x11-filter ssh-shell exec port-forward
    next
end

```

Oversized files

A couple of variations on the example could have to do with the processing of oversized files at a level other than the default setting. The ways that it can be approached are:

Set a non default threshold size and block the files

Set a non default threshold size and not scan the files over the threshold but allow them to pass through the FortiGate firewall.

In the following instructions:

- We will just use 2 MB as the new threshold.
- In the CLI instructions we will limit the configuration to just the HTTP settings for the purposes of brevity and simplicity.

Option 1

Option 1 can be done in the GUI.

1. Go to **Policy & Objects > Policy > Proxy Options**.
2. Edit the profile that you want to change in.
3. Fill out the fields with the following information:

Common Options

Comfort Clients	enabled
------------------------	---------

Interval (Seconds)	15
Amount(bytes)	1
Block Oversized File/Email	enabled
Threshold(MB)	2

4. Select **OK**.

5. Enter the following CLI command:

```
config firewall profile-protocol-options
edit example_standard
config http
set options clientcomfort no-content-summary oversize
set oversize-limit 2
set comfort-interval 15
next
end
```

Option 2

Option 2 can only be done in the CLI. Enter the following CLI command:

```
config firewall profile-protocol-options
edit example_standard
config http
set options clientcomfort no-content-summary
set oversize-limit 2
set comfort-interval 15
next
end
```

Example: DoS policy

The company wishes to protect against Denial of Service attack. They have chosen some where they wish to block the attacks of the incidence goes above a certain threshold and for some others they are just trying to get a baseline of activity for those types of attacks so they are letting the traffic pass through without action.

- The interface to the Internet is on WAN1
- There is no requirement to specify which addresses are being protected or protected from.
- The protection is to extend to all services.
- The TCP attacks are to be blocked
- The UDP, ICMP, and IP attacks are to be recorded but not blocked.
- The tcp_syn_flood attack's threshold is to be changed from the default to 1000

1. Go to **Policy & Objects > Policy > DoS**.
2. Create a new policy
3. Fill out the fields with the following information:

Incoming Interface	wan1
Source Address	all
Destination Addresses	all
Service	ALL

Anomalies

Name	Status	Logging	Action	Threshold
tcp_syn_flood	enabled	enabled	Block	1000
tcp_port_scan	enabled	enabled	Block	<default value>
tcp_src_session	enabled	enabled	Block	<default value>
tcp_dst_session	enabled	enabled	Block	<default value>
udp_flood	enabled	enabled	Pass	<default value>
udp_scan	enabled	enabled	Pass	<default value>
udp_src_session	enabled	enabled	Pass	<default value>
udp_dst_session	enabled	enabled	Pass	<default value>
icmp_flood	enabled	enabled	Pass	<default value>
icmp_sweep	enabled	enabled	Pass	<default value>
icmp_src_session	enabled	enabled	Pass	<default value>
icmp_dst_session	enabled	enabled	Pass	<default value>
ip_src_session	enabled	enabled	Pass	<default value>
ip_dst_session	enabled	enabled	Pass	<default value>
sctp_flood	not enabled	not enabled	Pass	<default value>
sctp_scan	not enabled	not enabled	Pass	<default value>
sctp_src_session	not enabled	not enabled	Pass	<default value>
sctp_dst_session	not enabled	not enabled	Pass	<default value>

4. Toggle the button next to **Enable this policy** to **ON**.
5. Select **OK**.
6. Enter the following CLI command:


```
config firewall DoS-policy
edit 0
    set status enable
    set interface ''
    config anomaly
        edit "tcp_syn_flood"
            set status enable
            set log enable
            set action block
            set threshold 1000
        next
        edit "tcp_port_scan"
            set status enable
            set log enable
            set action block
        next
        edit "tcp_src_session"
            set status enable
            set log enable
            set action block
        next
        edit "tcp_dst_session"
            set status enable
            set log enable
            set action block
    next
    edit "udp_flood"
        set status enable
        set log enable
    next
    edit "udp_scan"
        set status disable
        set status enable
        set log enable
    next
    edit "udp_src_session"
        set status enable
        set log enable
    next
    edit "udp_dst_session"
        set status enable
        set log enable
    next
    edit "icmp_flood"
        set status enable
        set log enable
    next
    edit "icmp_sweep"
        set status enable
        set log enable
    next
    edit "icmp_src_session"
        set status enable
        set log enable
    next
    edit "icmp_dst_session"
        set status enable
```

```
        set log enable
        next
    edit "ip_src_session"
        set status enable
        set log enable
        next
    edit "ip_dst_session"
        set status enable
        set log enable
        next
    end
next
end
```

Multicast forwarding

Multicasting (also called IP multicasting) consists of using a single multicast source to send data to many receivers. Multicasting can be used to send data to many receivers simultaneously while conserving bandwidth and reducing network traffic. Multicasting can be used for one-way delivery of media streams to multiple receivers and for one-way data transmission for news feeds, financial information, and so on.

Also RIPv2 uses multicasting to share routing table information, OSPF uses multicasting to send hello packets and routing updates, Enhanced Interior Gateway Routing Protocol (EIGRP) uses multicasting to send routing information to all EIGRP routers on a network segment and the Bonjour network service uses multicasting for DNS.

A FortiGate unit can operate as a Protocol Independent Multicast (PIM) version 2 router. FortiGate units support PIM sparse mode (RFC 4601) and PIM dense mode (RFC 3973) and can service multicast servers or receivers on the network segment to which a FortiGate unit interface is connected. Multicast routing is not supported in transparent mode (TP mode).



To support PIM communications, the sending/receiving applications and all connecting PIM routers in between must be enabled with PIM version 2. PIM can use static routes, RIP, OSPF, or BGP to forward multicast packets to their destinations. To enable source-to-destination packet delivery, either sparse mode or dense mode must be enabled on the PIM-router interfaces. Sparse mode routers cannot send multicast messages to dense mode routers. In addition, if a FortiGate unit is located between a source and a PIM router, two PIM routers, or is connected directly to a receiver, you must create a security policy manually to pass encapsulated (multicast) packets or decapsulated data (IP traffic) between the source and destination.

A PIM domain is a logical area comprising a number of contiguous networks. The domain contains at least one Bootstrap Router (BSR), and if sparse mode is enabled, a number of Rendezvous Points (RPs) and Designated Routers (DRs). When PIM is enabled on a FortiGate unit, the FortiGate unit can perform any of these functions at any time as configured.

Sparse mode

Initially, all candidate BSRs in a PIM domain exchange bootstrap messages to select one BSR to which each RP sends the multicast address or addresses of the multicast group(s) that it can service. The selected BSR chooses one RP per multicast group and makes this information available to all of the PIM routers in the domain through bootstrap messages. PIM routers use the information to build packet distribution trees, which map each multicast group to a specific RP. Packet distribution trees may also contain information about the sources and receivers associated with particular multicast groups.



When a FortiGate unit interface is configured as a multicast interface, sparse mode is enabled on it by default to ensure that distribution trees are not built unless at least one downstream receiver requests multicast traffic from a specific source. If the sources of multicast traffic and their receivers are close to each other and the PIM domain contains a dense population of active receivers, you may choose to enable dense mode throughout the PIM domain instead.

An RP represents the root of a non-source-specific distribution tree to a multicast group. By joining and pruning the information contained in distribution trees, a single stream of multicast packets (for example, a video feed) originating from the source can be forwarded to a certain RP to reach a multicast destination.

Each PIM router maintains a Multicast Routing Information Base (MRIB) that determines to which neighboring PIM router join and prune messages are sent. An MRIB contains reverse-path information that reveals the path of a multicast packet from its source to the PIM router that maintains the MRIB.

To send multicast traffic, a server application sends IP traffic to a multicast group address. The locally elected DR registers the sender with the RP that is associated with the target multicast group. The RP uses its MRIB to forward a single stream of IP packets from the source to the members of the multicast group. The IP packets are replicated only when necessary to distribute the data to branches of the RP's distribution tree.

To receive multicast traffic, a client application can use Internet Group Management Protocol (IGMP) version 1 (RFC 1112), 2 (RFC 2236), or 3 (RFC 3376) control messages to request the traffic for a particular multicast group. The locally elected DR receives the request and adds the host to the multicast group that is associated with the connected network segment by sending a join message towards the RP for the group. Afterward, the DR queries the hosts on the connected network segment continually to determine whether the hosts are active. When the DR no longer receives confirmation that at least one member of the multicast group is still active, the DR sends a prune message towards the RP for the group.

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with RFC 4601: Protocol Independent Multicast - Sparse Mode (PIM-SM). You can use the following command to configure IPv6 PIM sparse multicast routing.

```
config router multicast6
    set multicast-routing {enable | disable}
    config interface
        edit <interface-name>
            set hello-interval <1-65535 seconds>
            set hello-holdtime <1-65535 seconds>
        end
    config pim-sm-global
        config rp-address
            edit <index>
                set ipv6-address <ipv6-address>
            end
        end
    end
```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```
diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute
```

Dense mode

The packet organization used in sparse mode is also used in dense mode. When a multicast source begins to send IP traffic and dense mode is enabled, the closest PIM router registers the IP traffic from the multicast source (S) and forwards multicast packets to the multicast group address (G). All PIM routers initially broadcast the multicast packets throughout the PIM domain to ensure that all receivers that have requested traffic for multicast group address G can access the information if needed.

To forward multicast packets to specific destinations afterward, the PIM routers build distribution trees based on the information in multicast packets. Upstream PIM routers depend on prune/graft messages from downstream PIM routers to determine if receivers are actually present on directly connected network segments. The PIM routers exchange state refresh messages to update their distribution trees. FortiGate units store this state information in a Tree Information Base (TIB), which is used to build a multicast forwarding table. The information in the multicast forwarding table determines whether packets are forwarded downstream. The forwarding table is updated whenever the TIB is modified.

PIM routers receive data streams every few minutes and update their forwarding tables using the source (S) and multicast group (G) information in the data stream. Superfluous multicast traffic is stopped by PIM routers that do not have downstream receivers—PIM routers that do not manage multicast groups send prune messages to the upstream PIM routers. When a receiver requests traffic for multicast address G, the closest PIM router sends a graft message upstream to begin receiving multicast packets.

FortiGate units operating in NAT mode can also be configured as multicast routers. You can configure a FortiGate unit to be a Protocol Independent Multicast (PIM) router operating in Sparse Mode (SM) or Dense Mode (DM).

Multicast IP addresses

Multicast uses the Class D address space. The 224.0.0.0 to 239.255.255.255 IP address range is reserved for multicast groups. The multicast address range applies to multicast groups, not to the originators of multicast packets. The following table lists the reserved multicast address ranges and describes what they are reserved for:

Reserved Multicast address ranges

Reserved Address Range	Use	Notes
224.0.0.0 to 224.0.0.255	Used for network protocols on local networks. For more information, see RFC 1700.	In this range, packets are not forwarded by the router but remain on the local network. They have a Time to Live (TTL) of 1. These addresses are used for communicating routing information.
224.0.1.0 to 238.255.255.255	Global addresses used for multicasting data between organizations and across the Internet. For more information, see RFC 1700.	Some of these addresses are reserved, for example, 224.0.1.1 is used for Network Time Protocol (NTP).

Reserved Address Range	Use	Notes
239.0.0.0 to 239.255.255.255	Limited scope addresses used for local groups and organizations. For more information, see RFC 2365.	Routers are configured with filters to prevent multicasts to these addresses from leaving the local system.

Creating multicast security policies requires multicast firewall addresses. You can add multicast firewall addresses by going to **Firewall Objects > Address > Addresses** and selecting **Create New > Multicast Address**. The factory default configuration includes multicast addresses for Bonjour (224.0.0.251-224.0.0.251), EIGRP (224.0.0.10-224.0.0.100), OSPF (224.0.0.5-224.0.0.60), all_hosts (224.0.0.1-224.0.0.1), and all_routers (224.0.0.2-224.0.0.2).

PIM Support

A FortiGate unit can be configured to support PIM by going to **Router > Dynamic > Multicast** and enabling multicast routing. You can also enable multicast routing using the `config router multicast` CLI command. When PIM is enabled, the FortiGate unit allocates memory to manage mapping information. The FortiGate unit communicates with neighboring PIM routers to acquire mapping information and if required, processes the multicast traffic associated with specific multicast groups.



The end-user multicast client-server applications must be installed and configured to initiate Internet connections and handle broadband content such as audio/video information.

Client applications send multicast data by registering IP traffic with a PIM-enabled router. An end-user could type in a class D multicast group address, an alias for the multicast group address, or a call-conference number to initiate the session.

Rather than sending multiple copies of generated IP traffic to more than one specific IP destination address, PIM-enabled routers encapsulate the data and use the one multicast group address to forward multicast packets to multiple destinations. Because one destination address is used, a single stream of data can be sent. Client applications receive multicast data by requesting that the traffic destined for a certain multicast group address be delivered to them — end-users may use phone books, a menu of ongoing or future sessions, or some other method through a user interface to select the address of interest.

A class D address in the 224.0.0.0 to 239.255.255.255 range may be used as a multicast group address, subject to the rules assigned by the Internet Assigned Numbers Authority (IANA). All class D addresses must be assigned in advance. Because there is no way to determine in advance if a certain multicast group address is in use, collisions may occur (to resolve this problem, end-users may switch to a different multicast address).

To configure a PIM domain

1. If you will be using sparse mode, determine appropriate paths for multicast packets.
2. Make a note of the interfaces that will be PIM-enabled. These interfaces may run a unicast routing protocol.
3. If you will be using sparse mode and want multicast packets to be handled by specific (static) RPs, record the IP addresses of the PIM-enabled interfaces on those RPs.

4. Enable PIM version 2 on all participating routers between the source and receivers. On FortiGate units, use the `config router multicast` command to set global operating parameters.
5. Configure the PIM routers that have good connections throughout the PIM domain to be candidate BSRs.
6. If sparse mode is enabled, configure one or more of the PIM routers to be candidate RPs.
7. If required, adjust the default settings of PIM-enabled interface(s).

Multicast forwarding and FortiGate units

In both transparent mode and NAT mode you can configure FortiGate units to forward multicast traffic.

For a FortiGate unit to forward multicast traffic you must add FortiGate multicast security policies. Basic multicast security policies accept any multicast packets at one FortiGate interface and forward the packets out another FortiGate interface. You can also use multicast security policies to be selective about the multicast traffic that is accepted based on source and destination address, and to perform NAT on multicast packets.

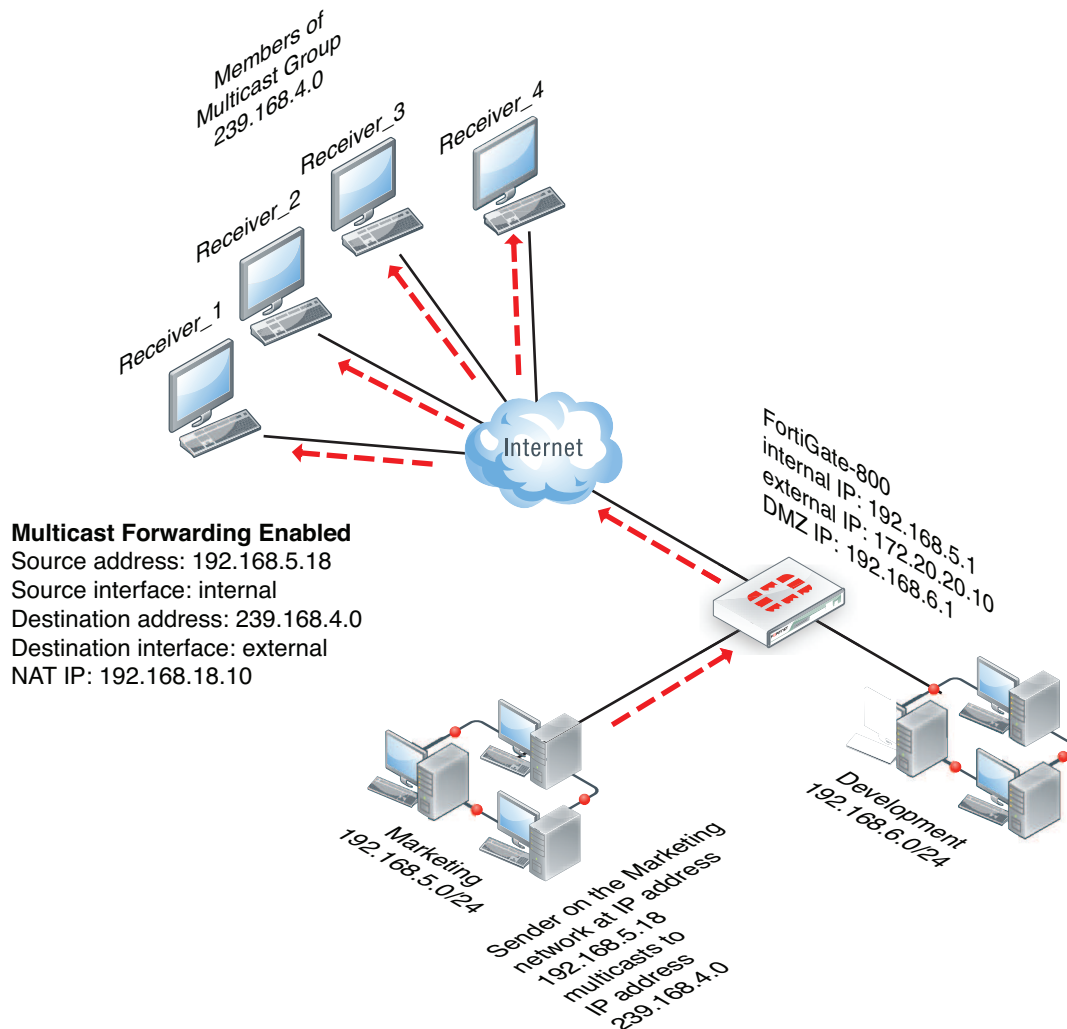
In the example shown below, a multicast source on the Marketing network with IP address 192.168.5.18 sends multicast packets to the members of network 239.168.4.0. At the FortiGate unit, the source IP address for multicast packets originating from workstation 192.168.5.18 is translated to 192.168.18.10. In this example, the FortiGate unit is not acting as a multicast router.

Multicast forwarding and RIPv2

RIPv2 uses multicast to share routing table information. If your FortiGate unit is installed on a network that includes RIPv2 routers, you must configure the FortiGate unit to forward multicast packets so that RIPv2 devices can share routing data through the FortiGate unit. No special FortiGate configuration is required to share RIPv2 data, you can simply use the information in the following sections to configure the FortiGate unit to forward multicast packets.



RIPv1 uses broadcasting to share routing table information. To allow RIPv1 packets through a FortiGate unit you can add standard security policies. Security policies to accept RIPv1 packets can use the ANY predefined firewall service or the RIP predefined firewall service.

Example multicast network including a FortiGate unit that forwards multicast packets

Configuring FortiGate multicast forwarding

You configure FortiGate multicast forwarding from the Command Line Interface (CLI). Two steps are required:

- Adding multicast security policies
- Enabling multicast forwarding

This second step is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.



There is sometimes a confusion between the terms “forwarding” and “routing”. These two functions should not be taking place at the same time.

It is mentioned that multicast-forward should be enabled when the FortiGate unit is in NAT mode and that this will forward any multicast packet to all interfaces. However, this parameter should **NOT** be enabled when the FortiGate unit operates as a multicast router (i.e. with a routing protocol enabled). It should only be enabled when there is no routing protocols activated.

Adding multicast security policies

You need to add security policies to allow packets to pass from one interface to another. Multicast packets require multicast security policies. You add multicast security policies from the CLI using the `config firewall multicast-policy` command. As with unicast security policies, you specify the source and destination interfaces and optionally the allowed address ranges for the source and destination addresses of the packets.

You can also use multicast security policies to configure source NAT and destination NAT for multicast packets.

Keep the following in mind when configuring multicast security policies:

- The matched forwarded (outgoing) IP multicast source IP address is changed to the configured IP address.
- Source and Destination interfaces are optional. If left blank, then the multicast will be forwarded to ALL interfaces.
- Source and Destination addresses are optional. If left un set, then it will mean ALL addresses.
- The `nat` keyword is optional. Use it when source address translation is needed.

Enabling multicast forwarding

Multicast forwarding is enabled by default. In NAT mode you must use the `multicast-forward` keyword of the `system settings` CLI command to enable or disable multicast forwarding. When `multicast-forward` is enabled, the FortiGate unit forwards any multicast IP packets in which the TTL is 2 or higher to all interfaces and VLAN interfaces except the receiving interface. The TTL in the IP header will be reduced by 1. Even though the multicast packets are forwarded to all interfaces, you must add security policies to actually allow multicast packets through the FortiGate. In our example, the security policy allows multicast packets received by the internal interface to exit to the external interface.



Enabling multicast forwarding is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

Enter the following CLI command to enable multicast forwarding:

```
config system settings
  set multicast-forward enable
end
```

If multicast forwarding is disabled and the FortiGate unit drops packets that have multicast source or destination addresses.

You can also use the `multicast-ttl-notchange` keyword of the `system settings` command so that the FortiGate unit does not increase the TTL value for forwarded multicast packets. You should use this option only if packets are expiring before reaching the multicast router.

```
config system settings
    set multicast-ttl-notchange enable
end
```

In transparent mode, the FortiGate unit does not forward frames with multicast destination addresses. Multicast traffic such as the one used by routing protocols or streaming media may need to traverse the FortiGate unit, and should not be interfere with the communication. To avoid any issues during transmission, you can set up multicast security policies. These types of security policies can only be enabled using the CLI.



The CLI parameter `multicast-skip-policy` must be disabled when using multicast security policies. To disable enter the command

```
config system settings
    set multicast-skip-policy disable
end
```

In this simple example, no check is performed on the source or destination interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces on the forwarding domain, except the incoming interface.

To enable the multicast policy

```
config firewall multicast-policy
    edit 1
        set action accept
    end
```

In this example, the multicast policy only applies to the source port of WAN1 and the destination port of Internal.

To enable the restrictive multicast policy

```
config firewall multicast-policy
    edit 1
        set srcintf wan1
        set dstintf internal
        set action accept
    end
```

In this example, packets are allowed to flow from WAN1 to Internal, and sourced by the address 172.20.120.129, which is represented by the address object "example_addr-1".

To enable the restrictive multicast policy

```
config firewall multicast-policy
    edit 1
        set srcintf wan1
        set srcaddr example_addr-1
        set dstintf internal
        set action accept
    end
```

This example shows how to configure the multicast security policy required for the configuration shown. This policy accepts multicast packets that are sent from a PC with IP address 192.168.5.18 to destination address range 239.168.4.0. The policy allows the multicast packets to enter the internal interface and then exit the external interface. When the packets leave the external interface their source address is translated to 192.168.18.10

```
config firewall multicast-policy
```

```
edit 5
  set srcaddr 192.168.5.18 255.255.255.255
  set srcintf internal
  set destaddr 239.168.4.0 255.255.255.0
  set dstintf external
  set nat 192.168.18.10
end
```

This example shows how to configure a multicast security policy so that the FortiGate unit forwards multicast packets from a multicast Server with an IP 10.10.10.10 is broadcasting to address 225.1.1.1. This Server is on the network connected to the FortiGate DMZ interface.

```
config firewall multicast-policy
edit 1
    set srcintf DMZ
    set srcaddr 10.10.10.10 255.255.255.255
    set dstintf Internal
    set dstaddr 225.1.1.1 255.255.255.255
    set action accept
edit 2
    set action deny
end
```

Multicast routing examples

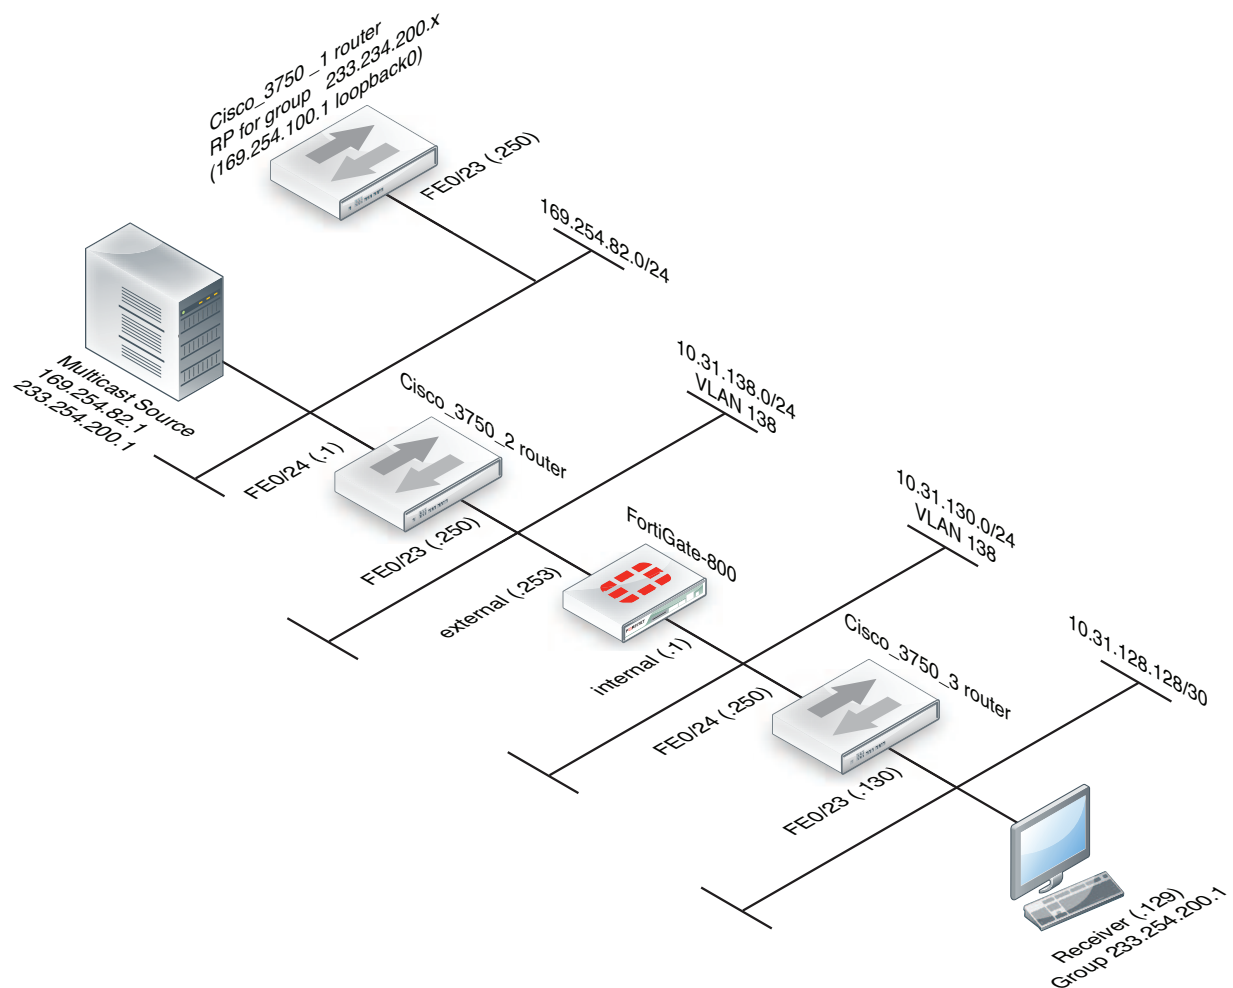
This section contains the following multicast routing configuration examples and information:

- Example FortiGate PIM-SM configuration using a static RP
- FortiGate PIM-SM debugging examples
- Example multicast destination NAT (DNAT) configuration
- Example PIM configuration that uses BSR to find the RP

Example FortiGate PIM-SM configuration using a static RP

The example Protocol Independent Multicast Sparse Mode (PIM-SM) configuration shown below has been tested for multicast interoperability using PIM-SM between Cisco 3750 switches running 12.2 and a FortiGate-800 running FortiOS v3.0 MR5 patch 1. In this configuration, the receiver receives the multicast stream when it joins the group 233.254.200.1.

Example FortiGate PIM-SM topology



The configuration uses a statically configured rendezvous point (RP) which resides on the Cisco_3750_1. Using a bootstrap router (BSR) was not tested in this example. See “Example PIM configuration that uses BSR to find the RP” for an example that uses a BSR.

Configuration steps

The following procedures show how to configure the multicast configuration settings for the devices in the example configuration.

- Cisco_3750_1 router configuration
- Cisco_3750_2 router configuration
- To configure the FortiGate-800 unit
- Cisco_3750_3 router configuration

Cisco_3750_1 router configuration

```
version 12.2
!
hostname Cisco-3750-1
```

```

!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface Loopback0
    ip address 169.254.100.1 255.255.255.255
!
interface FastEthernet1/0/23
    switchport access vlan 182
    switchport mode access
!
interface FastEthernet1/0/24
    switchport access vlan 172
    switchport mode access
!
interface Vlan172
    ip address 10.31.138.1 255.255.255.0
    ip pim sparse-mode
    ip igmp query-interval 125
    ip mroute-cache distributed
!
interface Vlan182
    ip address 169.254.82.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 169.254.82.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
ip access-list standard Source-RP
    permit 233.254.200.0 0.0.0.255

```

Cisco_3750_2 router configuration

```

version 12.2
!
hostname Cisco-3750-2
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!

```

```

interface FastEthernet1/0/23
    switchport access vlan 138
    switchport mode access
!
interface FastEthernet1/0/24
    switchport access vlan 182
    switchport mode access
!
interface Vlan138
    ip address 10.31.138.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
interface Vlan182
    ip address 169.254.82.1 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.138.253
ip route 169.254.100.1 255.255.255.255 169.254.82.250
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255

```

To configure the FortiGate-800 unit

1. Configure the internal and external interfaces.

- **Internal**

Go to **System > Network > Interfaces**.

Select the internal interface.

Verify the following settings:

Type:	Physical Interface
Addressing mode:	Manual
IP/Network Mask:	10.31.138.253 255.255.255.0
Administrative Access:	PING

Select **OK**.

- **External**

Go to **System > Network > Interfaces**.

Select the external interface.

Verify the following settings:

Type:	Physical Interface
Addressing mode:	Manual
IP/Network Mask:	10.31.130.253 255.255.255.0
Administrative Access:	HTTPS and PING

Select **OK**.

2. Add a firewall addresses.

Go to **Policy & Objects > Objects > Addresses**.

- RP

Select **Create New**.

Use the following settings:

Category:	Address
Name:	RP
Type:	Subnet
Subnet/IP Range:	169.254.100.1/32
Interface:	Any
Visibility:	<enabled>

Select **OK**.

- Multicast source subnet

Select **Create New**.

Use the following settings:

Category:	Address
Name:	multicast_source_subnet
Type:	Subnet
Subnet/IP Range:	169.254.82.0/24
Interface:	Any
Visibility:	<enabled>

Select **OK**.

3. Add destination multicast address

Go to **Policy & Objects > Objects > Addresses**.

Select **Create New**.

Use the following settings:

Category:	Multicast Address
Name:	Multicast_stream
Type:	Broadcast Subnet
Broadcast Subnet:	233.254.200.0/24
Interface:	Any
Visibility:	<enabled>

Select **OK**.

4. Add standard security policies to allow traffic to reach the RP.

Go to **Policy & Objects > Policy > IPv4**.

- 1st policy

Select **Create New**

Use the following settings:

Incoming Interface:	internal
Source Address:	all
Outgoing Interface:	external
Destination Address:	RP
Schedule:	always
Service:	ALL
Action:	ACCEPT

Select **OK**.

- 2nd policy

Select **Create New**

Use the following settings:

Incoming Interface:	external
Source Address:	RP

Outgoing Interface:	internal
Destination Address:	all
Schedule:	always
Service:	ALL
Action:	ACCEPT

Select **OK**.

5. Add the multicast security policy.

Go to **Policy & Objects > Policy > Multicast**.

Select **Create New**.

Use the following settings:

Incoming Interface:	external
Source Address:	multicast_source_subnet
Outgoing Interface:	internal
Destination Address:	multicast_stream
Protocol:	Any
Action:	ACCEPT

Select **OK**.

6. Add an access list. (CLI only)

```
config router access-list
edit Source-RP
config rule
edit 1
set prefix 233.254.200.0 255.255.255.0
set exact-match disable
next
end
```

7. Add some static routes.

Go to **Router > Static > Static Routes**.

- Route 1

Select **Create New**.

Use the following settings:

Destination IP/Mask:	0.0.0.0/0.0.0.0
-----------------------------	-----------------

Device:	internal
Gateway:	10.31.130.250
Distance:	<default>
Priority:	<default>

Select **OK**.

- Route 2

Select **Create New**.

Use the following settings:

Destination IP/Mask:	169.254.0.0/16
Device:	external
Gateway:	10.31.138.250
Distance:	<default>
Priority:	<default>

Select **OK**.

8. Configure multicast routing.

Go to **Router > Dynamic > Multicast**.

Add the following Static Rendezvous Point(s):

- 169.254.100.1
- Route 1

Select **Create New**.

Use the following settings:

Interface:	internal
PIM Mode:	Sparse Mode
DR Priority:	<not needed in this scenario>
RP Candidate:	<not needed in this scenario>
RP Candidate Priority:	<not needed in this scenario>

Select **OK**.

- Route 2

Select **Create New**.

Use the following settings:

Interface:	external
PIM Mode:	Sparse Mode
DR Priority:	
RP Candidate:	
RP Candidate Priority:	

Select **OK**.

Cisco_3750_3 router configuration

```

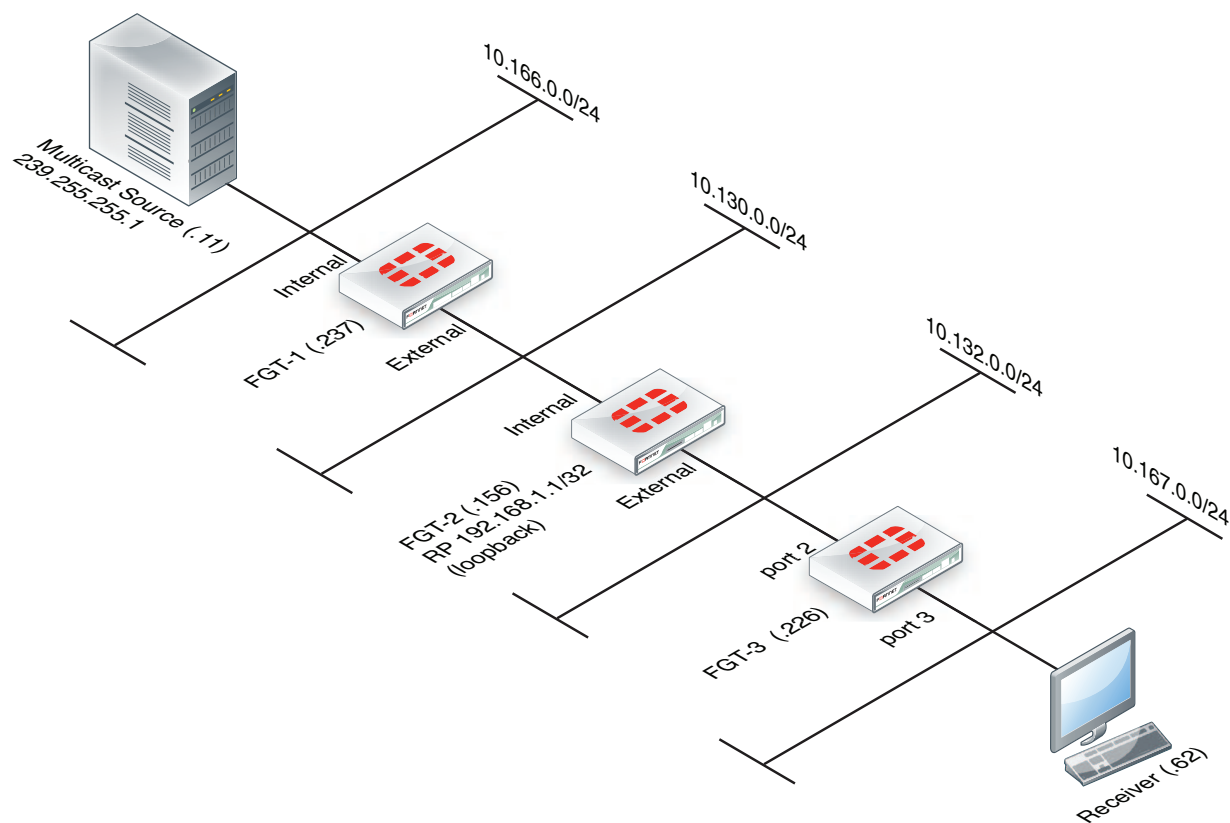
version 12.2
!
hostname Cisco-3750-3
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
    switchport access vlan 128
    switchport mode access
!
interface FastEthernet1/0/24
    switchport access vlan 130
    switchport mode access
!
interface Vlan128
    ip address 10.31.128.130 255.255.255.252
    ip pim sparse-mode
    ip mroute-cache distributed
!
interface Vlan130
    ip address 10.31.130.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.130.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255

```

FortiGate PIM-SM debugging examples

Using the example topology shown below, you can trace the multicast streams and states within the three FortiGate units (FGT-1, FGT-2, and FGT-3) using the debug commands described in this section. The command output in this section is taken from FortiGate unit when the multicast stream is flowing correctly from source to receiver.

PIM-SM debugging topology



Checking that the receiver has joined the required group

From the last hop router, FGT-3, you can use the following command to check that the receiver has correctly joined the required group.

```
FGT-3 # get router info multicast igmp groups
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
239.255.255.1 port3 00:31:15 00:04:02 10.167.0.62
```

Only 1 receiver is displayed for a particular group, this is the device that responded to the IGMP query request from the FGT-3. If a receiver is active the expire time should drop to approximately 2 minutes before being refreshed.

Checking the PIM-SM neighbors

Next the PIM-SM neighbors should be checked. A PIM router becomes a neighbor when the PIM router receives a PIM hello. Use the following command to display the PIM-SM neighbors of FGT-3.

```
FGT-3 # get router info multicast pim sparse-mode neighbour
Neighbor Interface Uptime/Expires Ver DR
Address Priority/Mode
10.132.0.156 port2 01:57:12/00:01:33 v2 1 /
```

Checking that the PIM router can reach the RP

The rendezvous point (RP) must be reachable for the PIM router (FGT-3) to be able to send the *, G join to request the stream. This can be checked for FGT-3 using the following command:

```
FGT-3 # get router info multicast pim sparse-mode rp-mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
RP: 192.168.1.1
Uptime: 07:23:00
```

Viewing the multicast routing table (FGT-3)

The FGT-3 unicast routing table can be used to determine the path taken to reach the RP at 192.168.1.1. You can then check the stream state entries using the following commands:

```
FGT-3 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

(*,*,RP) Entries	This state may be reached by general joins for all groups served by a specified RP.
(*,G) Entries	State that maintains the RP tree for a given group.
(S,G) Entries	State that maintains a source-specific tree for source S and group G.
(S,G,rpt) Entries	State that maintains source-specific information about source s on the RP tree for G. For example, if a source is being received on the source-specific tree, it will normally have been pruned off the RP tree.
FCR	The FCR state entries are for tracking the sources in the <*, G> when <S, G> is not available for any reason, the stream would typically be flowing when this state exists.

Breaking down each entry in detail:

```
(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: JOINED
Local:
port3
Joined:
Asserted:
FCR:
```

The RP will always be listed in a *, G entry, the RPF neighbor and interface index will also be shown. In this topology these are the same in all downstream PIM routers. The state is active so the upstream state is joined.

In this case FGT-3 is the last hop router so the IGMP join is received locally on port3. There is no PIM outgoing interface listed for this entry as it is used for the upstream PIM join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.132.0.156
RPF idx: port2
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
Asserted:
Outgoing:
port3
```

This is the entry for the SPT, no RP IS listed. The S, G stream will be forwarded out of the stated outgoing interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: NOT PRUNED
Local:
Pruned:
Outgoing:
```

The above S, G, RPT state is created for all streams that have both a S, G and a *, G entry on the router. This is not pruned in this case because of the topology, the RP and source are reachable over the same interface.

Although not seen in this scenario, assert states may be seen when multiple PIM routers exist on the same LAN which can lead to more than one upstream router having a valid forwarding state. Assert messages are used to elect a single forwarder from the upstream devices.

Viewing the PIM next-hop table

The PIM next-hop table is also very useful for checking the various states, it can be used to quickly identify the states of multiple multicast streams

```
FGT-3 # get router info multicast pim sparse-mode next-hop
Flags: N = New, R = RP, S = Source, U = Unreachable
Destination Type Nexthop Nexthop Nexthop Metric Pref Refcnt
Num Addr Ifindex
-----
10.166.0.11 ..S. 1 10.132.0.156 9 21 110 3
192.168.1.1 .R.. 1 10.132.0.156 9 111 110 2
```

Viewing the PIM multicast forwarding table

Also you can check the multicast forwarding table showing the ingress and egress ports of the multicast stream.

```
FGT-3 # get router info multicast table

IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface (TTL threshold)
```

```
(10.166.0.11, 239.255.255.1), uptime 04:02:55, stat expires 00:02:25
Owner PIM-SM, Flags: TF
Incoming interface: port2
Outgoing interface list:
port3 (TTL threshold 1)
```

Viewing the kernel forwarding table

Also the kernel forwarding table can be verified, however this should give similar information to the above command:

```
FGT-3 # diag ip multicast mroute
grp=239.255.255.1 src=10.166.0.11 intf=9 flags=(0x10000000)[ ] status=resolved
last_assert=2615136 bytes=1192116 pkt=14538 wrong_if=0 num_ifs=1
index(ttl)=[6(1),]
```

Viewing the multicast routing table (FGT-2)

If you check the output on FGT-2 there are some small differences:

```
FGT-2 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
Local:
Joined:
external
Asserted:
FCR:
```

The *, G entry now has a joined interface rather than local because it has received a PIM join from FGT-3 rather than a local IGMP join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.130.0.237
RPF idx: internal
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
external
Asserted:
Outgoing:
external
```

The S, G entry shows that we have received a join on the external interface and the stream is being forwarded out of this interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
```



```

RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: PRUNED
Local:
Pruned:
Outgoing:
External

```

The S, G, RPT is different from FGT-3 because FGT-2 is the RP, it has pruned back the SPT for the RP to the first hop router.

Viewing the multicast routing table (FGT-1)

FGT-1 again has some differences with regard to the PIM-SM states, there is no *, G entry because it is not in the path of a receiver and the RP.

```

FGT-1_master # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

```

Below the S, G is the SPT termination because this FortiGate unit is the first hop router, the RPF neighbor always shows as 0.0.0.0 because the source is local to this device. Both the joined and outgoing fields show as external because the PIM join and the stream is egressing on this interface.

```

(10.166.0.11, 239.255.255.1)
RPF nbr: 0.0.0.0
RPF idx: None
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
external
Asserted:
Outgoing:
external

```

The stream has been pruned back from the RP because the end-to-end SPT is flowing, there is no requirement for the stream to be sent to the RP in this case.

```

(10.166.0.11, 239.255.255.1, rpt)
RP: 0.0.0.0
RPF nbr: 10.130.0.156
RPF idx: external
Upstream State: RPT NOT JOINED
Local:
Pruned:
Outgoing:

```

Example multicast destination NAT (DNAT) configuration

The example topology shown and described below shows how to configure destination NAT (DNAT) for two multicast streams. Both of these streams originate from the same source IP address, which is 10.166.0.11. The example configuration keeps the streams separate by creating 2 multicast NAT policies.

In this example the FortiGate units have the following roles:

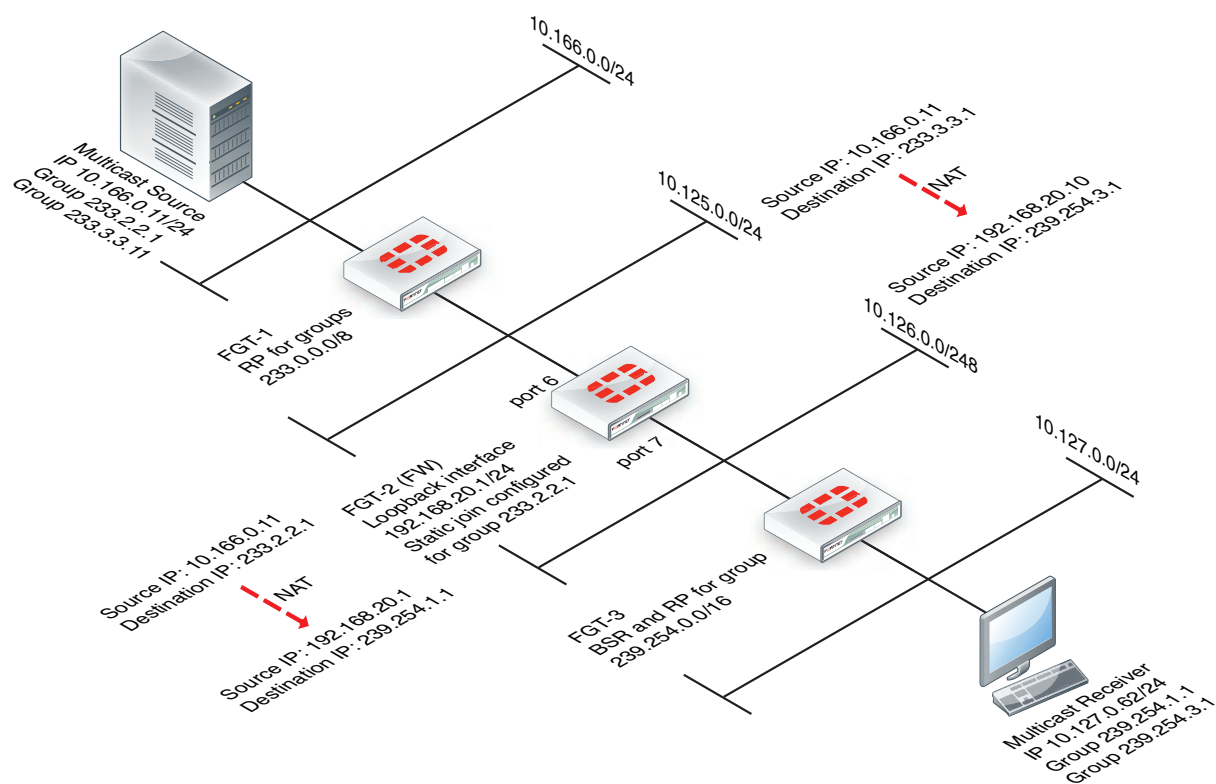
- FGT-1 is the RP for dirty networks, 233.0.0.0/8.
- FGT-2 performs all firewall and DNAT translations.
- FGT-3 is the RP for the clean networks, 239.254.0.0/16.
- FGT-1 and FGT-3 are functioning as PM enabled routers and could be replaced can be any PIM enabled router.

This example only describes the configuration of FGT-2.

FGT-2 performs NAT so that the receivers connected to FGT-3 receive the following translated multicast streams.

- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.2.2.1; FGT-3 translates the source and destination IPs to 192.168.20.1 and 239.254.1.1
- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.3.3.1; FGT-3 translates the source and destination IPs to 192.168.20.10 and 239.254.3.1

Example multicast DNAT topology



To configure FGT-2 for DNAT multicast

1. Add a loopback interface. In the example, the loopback interface is named `loopback`.

```
config system interface
  edit loopback
    set vdom root
    set ip 192.168.20.1 255.255.255.0
    set type loopback
  next
end
```

2. Add PIM and add a unicast routing protocol to the loopback interface as if it was a normal routed interface. Also add static joins to the loopback interface for any groups to be translated.

```
config router multicast
config interface
edit loopback
set pim-mode sparse-mode
config join-group
edit 233.2.2.1
next
edit 233.3.3.1
next
end
next
```

3. In this example, to add firewall multicast policies, different source IP addresses are required so you must first add an IP pool:

```
config firewall ippool
edit Multicast_source
set endip 192.168.20.20
set interface port6
set startip 192.168.20.10
next
end
```

4. Add the translation security policies.

Policy 2, which is the source NAT policy, uses the actual IP address of port6. Policy 1, the DNAT policy, uses an address from the IP pool. The source and destination addresses will need to be previously created address objects. For this example, 233.3.3.1 255.255.255.255 will be represented by "example-addr_1" and 10.166.0.11 255.255.255.255 will be represented by "example-addr_2". You will likely want to use something more intuitive from your own network.

```
config firewall multicast-policy
edit 1
set dnat 239.254.3.1
set dstaddr example-addr_1
set dstintf loopback
set nat 192.168.20.10
set srcaddr example-addr_2
set srcintf port6
next
edit 2
set dnat 239.254.1.1
set dstaddr 233.2.2.1 255.255.255.255
set dstintf loopback
set nat 192.168.20.1
set srcaddr 10.166.0.11 255.255.255.255
set srcintf port6
next
end
```

5. Add a firewall multicast policy to forward the stream from the loopback interface to the physical outbound interface.

This example is an any/any policy that makes sure traffic accepted by the other multicast policies can exit the FortiGate unit.

```

config firewall multicast-policy
edit 3
    set dstintf port7
    set srcintf loopback
next
end

```

Example PIM configuration that uses BSR to find the RP

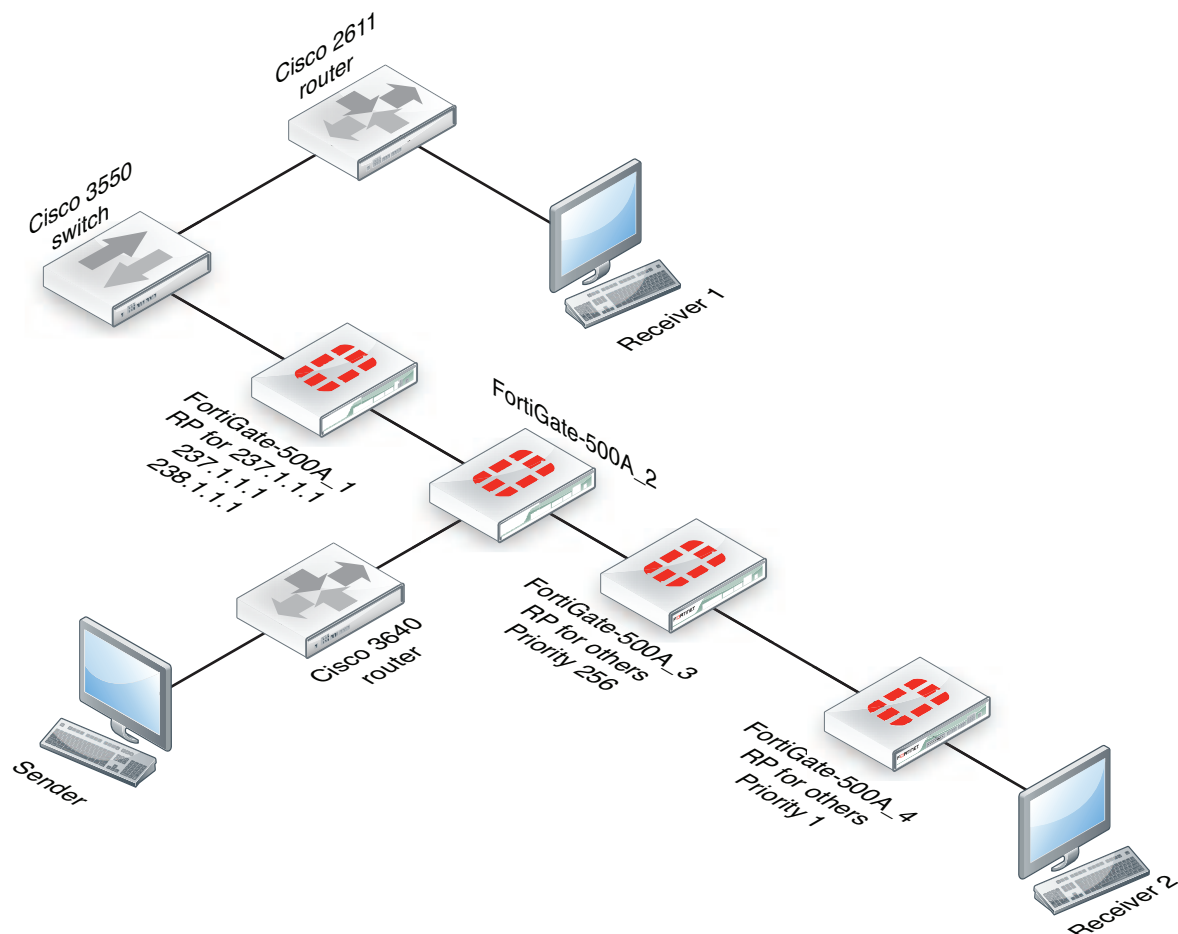
This example shows how to configure a multicast routing network for a network consisting of four FortiGate-500A units (FortiGate-500A_1 to FortiGate-500A_4). A multicast sender is connected to FortiGate-500A_2. FortiGate-500A_2 forwards multicast packets in two directions to reach Receiver 1 and Receiver 2.

The configuration uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface `lo0` must join the 236.1.1.1 group (source).

This example describes:

- Commands used in this example
- Configuration steps
- Example debug commands

PIM network topology using BSR to find the RP



Commands used in this example

This example uses CLI commands for the following configuration settings:

- Adding a loopback interface (lo0)
- Defining the multicast routing
- Adding the NAT multicast policy

Adding a loopback interface (lo0)

Where required, the following command is used to define a loopback interface named lo0.

```
config system interface
  edit lo0
    set vdom root
    set ip 1.4.50.4 255.255.255.255
    set allowaccess ping https ssh snmp http telnet
    set type loopback
  next
end
```

Defining the multicast routing

In this example, the following command syntax is used to define multicast routing.

The example uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface lo0 must join the 236.1.1.1 group (source).

```
config router multicast
  config interface
    edit port6
      set pim-mode sparse-mode
    next
    edit port1
      set pim-mode sparse-mode
    next
    edit lo0
      set pim-mode sparse-mode
      set rp-candidate enable
      config join-group
        edit 236.1.1.1
        next
      end
      set rp-candidate-priority 1
    next
  end
set multicast-routing enable
config pim-sm-global
  set bsr-allow-quick-refresh enable
  set bsr-candidate enable
  set bsr-interface lo0
  set bsr-priority 200
end
end
```

Adding the NAT multicast policy

In this example, the incoming multicast policy does the address translation.

The NAT address should be the same as the IP address of the of loopback interface. The DNAT address is the translated address, which should be a new group.

```
config firewall multicast-policy
edit 1
set dstintf port6
set srcintf lo0
next
edit 2
set dnat 238.1.1.1
set dstintf lo0
set nat 1.4.50.4
set srcintf port1
next
```

Configuration steps

In this sample, FortiGate-500A_1 is the RP for the group 228.1.1.1, 237.1.1.1, 238.1.1.1, and FortiGate-500A_4 is the RP for the other group which has a priority of 1. OSPF is used in this example to distribute routes including the loopback interface. All firewalls have full mesh security policies to allow any to any.

- In the FortiGate-500A_1 configuration, the NAT policy translates source address 236.1.1.1 to 237.1.1.1
- In the FortiGate-500A_4, configuration, the NAT policy translates source 236.1.1.1 to 238.1.1.1
- Source 236.1.1.1 is injected into network as well.

The following procedures include the CLI commands for configuring each of the FortiGate units in the example configuration.

To configure FortiGate-500A_1

1. Configure multicast routing.

```
config router multicast
config interface
edit port5
set pim-mode sparse-mode
next
edit port4
set pim-mode sparse-mode
next
edit lan
set pim-mode sparse-mode
next
edit port1
set pim-mode sparse-mode
next
edit lo999
set pim-mode sparse-mode
next
edit lo0
set pim-mode sparse-mode
set rp-candidate enable
set rp-candidate-group 1
next
end
set multicast-routing enable
config pim-sm-global
set bsr-candidate enable
```

```

        set bsr-interface lo0
    end
end

```

2. Add multicast security policies.

```

config firewall multicast-policy
    edit 1
        set dstintf port5
        set srcintf port4
    next
    edit 2
        set dstintf port4
        set srcintf port5
    next
    edit 3
    next
end

```

3. Add router access lists.

```

config router access-list
    edit 1
        config rule
            edit 1
                set prefix 228.1.1.1 255.255.255.255
                set exact-match enable
            next
            edit 2
                set prefix 237.1.1.1 255.255.255.255
                set exact-match enable
            next
            edit 3
                set prefix 238.1.1.1 255.255.255.255
                set exact-match enable
            next
        end
    next
end

```

To configure FortiGate-500A_2

1. Configure multicast routing.

```

config router multicast
    config interface
        edit "lan"
            set pim-mode sparse-mode
        next
    edit "port5"
        set pim-mode sparse-mode
    next
    edit "port2"
        set pim-mode sparse-mode
    next
    edit "port4"
        set pim-mode sparse-mode
    next
    edit "lo_5"

```

```
        set pim-mode sparse-mode
        config join-group
        edit 236.1.1.1
        next
    end
next
end
set multicast-routing enable
end
```

2. Add multicast security policies.

```
config firewall multicast-policy
edit 1
    set dstintf lan
    set srcintf port5
next
edit 2
    set dstintf port5
    set srcintf lan
next
edit 4
    set dstintf lan
    set srcintf port2
next
edit 5
    set dstintf port2
    set srcintf lan
next
edit 7
    set dstintf port1
    set srcintf port2
next
edit 8
    set dstintf port2
    set srcintf port1
next
edit 9
    set dstintf port5
    set srcintf port2
next
edit 10
    set dstintf port2
    set srcintf port5
next
edit 11
    set dnat 237.1.1.1
    set dstintf lo_5
    set nat 5.5.5.5
    set srcintf port2
next
edit 12
    set dstintf lan
    set srcintf lo_5
next
edit 13
    set dstintf port1
    set srcintf lo_5
```



```
next
edit 14
    set dstintf port5
    set srcintf lo_5
next
edit 15
    set dstintf port2
    set srcintf lo_5
next
edit 16
next
end
```

To configure FortiGate-500A_3

1. Configure multicast routing.

```
config router multicast
config interface
edit port5
    set pim-mode sparse-mode
next
edit port6
    set pim-mode sparse-mode
next
edit lo0
    set pim-mode sparse-mode
    set rp-candidate enable
    set rp-candidate-priority 255
next
edit lan
    set pim-mode sparse-mode
next
end
set multicast-routing enable
config pim-sm-global
    set bsr-candidate enable
    set bsr-interface lo0
end
end
```

2. Add multicast security policies.

```
config firewall multicast-policy
edit 1
    set dstintf port5
    set srcintf port6
next
edit 2
    set dstintf port6
    set srcintf port5
next
edit 3
    set dstintf port6
    set srcintf lan
next
edit 4
    set dstintf lan
    set srcintf port6
```

```
next
edit 5
    set dstintf port5
    set srcintf lan
next
edit 6
    set dstintf lan
    set srcintf port5
next
end
```

To configure FortiGate-500A_4

1. Configure multicast routing.

```
config router multicast
config interface
    edit port6
        set pim-mode sparse-mode
    next
    edit lan
        set pim-mode sparse-mode
    next
    edit port1
        set pim-mode sparse-mode
    next
    edit lo0
        set pim-mode sparse-mode
        set rp-candidate enable
        config join-group
            edit 236.1.1.1
            next
        end
        set rp-candidate-priority 1
    next
end
set multicast-routing enable
config pim-sm-global
set bsr-allow-quick-refresh enable
set bsr-candidate enable
set bsr-interface lo0
set bsr-priority 1
end
end
```

2. Add multicast security policies.

```
config firewall policy
    edit 1
        set srcintf lan
        set dstintf port6
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 2
        set srcintf port6
```

```
    set dstintf lan
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 3
    set srcintf port1
    set dstintf port6
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 4
    set srcintf port6
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 5
    set srcintf port1
    set dstintf lan
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 6
    set srcintf lan
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 7
    set srcintf port1
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 8
    set srcintf port6
    set dstintf lo0
    set srcaddr all
    set dstaddr all
```




High Performance Network Security



Copyright© 2017 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.